

usually only about \$1 or \$2 per megabyte and often is incurred only after the first Megabyte or two. Most prospective dial-up users will not find storage limitations or fees to be a major issue. Only if a user intends to regularly transfer files of one megabyte or larger, or intends to store files for long periods of time, will storage limitations or fees become a notable issue.

Pricing for multiple users. If Internet users in an organization will be using e-mail and participating in electronic discussion groups, it is almost imperative for each user to have his or her own Internet account. The pricing structures of dial-up Internet access providers accommodate multiple users from a single institution in different ways. Some make no special recognition in their fee structure, but others offer discounted annual or monthly base-level fees for additional accounts within an organization. And while some providers are pretty intransigent about staying with uniform "one user, one account" pricing, others will be willing to consider the specific circumstances of an organization that can bring a large number of accounts to the table and may be willing to work out something with institution-based group pricing levels.

Baud Rates, Units of Measurement for Connect Time, and SLIP/PPP

A final potential cost consideration concerns the relationship between communication speed and the method of measuring connect time. For those who use overwhelmingly e-mail and telnet and very little ftp, the implications of this relationship will probably be relatively minor. But for those who anticipate using ftp frequently to download large files to their own computers, this factor can be of considerable importance.

The transfer of a file from a remote host on the Internet to a user's computer is a two-step process in standard dial-up Internet access. The file first travels from the remote ftp site to the provider's computer. Most providers have a 56Kbps or higher connection to the Internet, and this is the speed at which the file travels between the remote ftp site and the provider's computer. The second step is the transfer of the file from the provider's computer to the user's computer.

The speed at which this happens depends on the baud rate at which the user is communicating with the provider's computer. For this user-to-provider connection, most Internet dial-up access providers can accommodate a variety of baud rates, but users typically have 9600 baud or lower modems. So while it may take only

a few minutes or even a matter of seconds for a file to travel between an ftp site and the provider's computer, it will take considerably longer for the file to travel the second leg of the journey from the provider's computer to the user.

For the Internet user who plans on frequently transferring large files via ftp, doing so at 9600 baud can seem interminably slow. In addition to the factor of time for its own sake, the baud speed at which files are to be downloaded also can have important cost implications. Normally, the unit of time measurement used by dial-up Internet providers is the time a user is connected to the provider's computers, not just to the Internet itself.

The time required for both segments of a dial-up file transfer — from the ftp site to the provider's computer and from the provider's computer to the user's — will likely be included as connect-time within a provider's pricing scheme. Thus, both parts will be fully subject to connect-hour charges under a straight connect-hour pricing scheme, or both will apply toward the monthly usage amount under a base-level pricing scheme. Under such a scenario, a large volume of file transfers at 9600 baud not only can be annoyingly slow, but expensive as well.

While 9600 baud is now a typical baud rate in dial-up Internet access, this limit can be stretched if the user has certain modem capabilities or software. On the latter score, some Internet providers offer what are known as SLIP (Serial Line Internet Protocol) or PPP (Point to Point Protocol) connections that allow a dial-up user to emulate a dedicated line user in many ways, including allowing for higher speed communications. This capability usually carries with it higher fees than standard dial-up. For the Internet user who will not be using ftp very frequently, the means to communicate at faster than 9600 baud will not be crucial. But if there will be a need to frequently transfer large files, the prospective user may want to consider what might be gained by selecting a provider offering a SLIP or PPP approach.

The cost elements described above should all be considered by the information specialist who is shopping among dial-up Internet access alternatives. Even when there is only one clearly convenient option and the comparative angle is not a consideration, it is important for the prospective user to understand the cost elements that can be important in approaching a provider and in planning for access to the Internet. But while costs are obviously important, they are not the only factors that should be considered in selecting or evaluating a provider, and the remainder of this article will take a look at some of those other factors.

Table 1. Sources of Low-Cost, Dial-Up Access to the Internet

1. University Computer Centers

The traditional avenue for those not affiliated with a large university that was a node on the Internet was to know someone who was. Many universities used to be fairly liberal about giving out courtesy passwords to people not affiliated with the university and usually at no charge. There is a clear trend away from this practice, however. More and more, universities are either limiting access strictly to their own community of scholars and students or are making guest accounts available to outside users for a fee.

2. Mid-level Regional Networks

The traditional service centers for the National Science Foundation's NSFnet that are so integral to the Internet are commonly referred to as regional mid-level networks. These networks have been responsible for installing and maintaining dedicated Internet connections in many institutions of higher education. There is considerable variation among these networks in the degree to which they are responsive to the nonacademic community and in the roles that they have assumed in providing low end dial-up Internet access. Thus, the extent to which a mid-level regional network is a realistic alternative for dial-up access depends on where one lives.

3. Nationwide Commercial Firms

A few companies in the commercial sector have seen electronic networking as a phenomenon with great market potential. The largest commercial provider of Internet access in the United States today is Performance Systems International, Inc. (PSI). PSI has done for commercial firms what the mid-level regional networks have done for academic institutions. There are additional vendors like CompuServe that offer varying levels of Internet access through more popular, low-priced dial-up services. Some offer the full complement of Internet functions — e-mail, telnet, and ftp — while others only offer e-mail.

4. Local Commercial Ventures

The rise in popularity of e-mail and bulletin boards has given rise to a proliferation of small-scale local commercial ventures that provide dial-up Internet access. Most efforts like these are located in major metropolitan areas, at least right now. Because of their limited budgets, most of these entrepreneurs spend very little on advertising or marketing and are sometimes hard to find out about. Many of them offer only e-mail, though some are beginning to provide telnet and ftp functionality.

5. Free-Nets

Community bulletin boards are being founded in an increasing number of cities in the United States. One of the earliest, the Cleveland Free-Net, serves as a model for a number of these efforts. Some community boards include, as one feature, a capability to connect to the Internet.

6. Library Networks and State Libraries

Many regional library networks in the U.S. are offering dial-in Internet access. A number of state libraries are also beginning to provide Internet access, and many more have it under consideration. As with most other options mentioned here, the extent to which these alternatives are available depends in large part on where a library is located.

System Performance Factors

The least expensive service is not the best solution if it does not meet the specific needs of a user. Conversely, there is no assurance that the most expensive service will be the most suitable, either. In the long run, costs will probably draw into closer correlation with quality of performance and service, but the business of providing Internet access is simply too new to yet assume so rational a pattern. At least for now, one should not anticipate that the cost of a service will necessarily pre-

dict of the quality of performance or support. Rather, one needs to examine performance and support criteria on their own merit. Several of the more important of these are discussed below.

Level of Functionality

Not all three of the main Internet functions — e-mail, telnet, and ftp — are available on all Internet access services. Many smaller local providers offer only e-mail functionality, as do some of the larger commercial providers. While more limited in function, these

providers typically are among the easiest with which to establish an account, so that a user needing only e-mail access may find them an attractive option. More and more, ftp and particularly telnet are seen as integral functions for libraries and information centers, and the prospective Internet user should be versed in what these functions have to offer before opting for only e-mail access.

Ease and Sophistication of Use

User interfaces in systems providing Internet access are not standardized, and there are many different e-mail systems and other components in use. Not surprisingly, there is great variation in the degree of user friendliness among these and in the breadth and sophistication of features available on them. In evaluating what Internet access providers have to offer, there are some user interface features that are usually indicative of a system's ease of use and sophistication.

Structure of basic user interface. The first system feature to look at is the basic user interface. There are still some Internet access providers who cater to an audience with a considerable amount of computer expertise, like a working knowledge of Unix. These systems are usually fully or nearly fully command-driven and are difficult to maneuver without some fairly high-level training. Providers with a more general user population in mind, on the other hand, have incorporated menus that make the task of getting out onto the Internet over their system fairly straightforward. As with any category of automated systems, some user interfaces are better designed than others. The best ones provide clear paths to select the functionality one wants to evoke — e-mail, telnet, or ftp — and enumerate options and supply prompts along the way.

Electronic mail component. In addition to the general user interface, another telltale feature of a system's ease of use is the e-mail component. This component usually consists of two parts: a mail manager that handles outgoing and incoming messages, and an editor on which one composes messages to send out — essentially a word processing system. A prospective user may notice a significant difference from one provider to another in the sophistication and ease of use of these e-mail managers and editors. Some editors, for example, have full-screen, very straightforward editing capabilities, while others have only line-by-line editing or require switching from one command mode to another in order to make changes during the course of typing a message. Different message handling software packages also vary in features like the sequence in which messages in a mail file are listed, the procedures for replying to messages, the availability of broadcast capabilities for sending a message simultaneously to more than one person, and so forth.

A prospective user may also want to look at e-mail components in terms of the procedures required for uploading and downloading messages. Particularly for longer and more formal communications, the user may not want to type a message "live" on the host system, but rather to compose it locally on disk and then upload it to send over the Internet. Similarly, many users also routinely download at least certain incoming mail messages to local disk. There can be quite a bit of variation in how different providers' systems work in this regard, and a prospective user may want to inquire about this in some detail.

Search aids. Another area that is useful for evaluating ease and sophistication of use is the extent to which various Internet search aids are available on a provider's system. Some of the more popular programs that assist users in locating and handling the myriad of resources available over the Internet include Gopher, WAIS, and Archie. Many Internet providers have these mounted on their hosts for direct access by subscribers. When they are not available on a provider's host, a user has to telnet to another site where they are mounted in order to use them. Thus, while it is usually still possible to make use of these aids, not having them available locally adds an inconvenient intermediate step.

The features described above — the general user interface, the e-mail system, and availability of Internet search aids on the host — by no means comprise an exhaustive list of the features that make one system easier to use or more sophisticated than another. But these three features do provide the prospective user with a frame of reference for comparing different Internet access providers. Armed with these three points of evaluation, the prospective subscriber will be able to get a pretty good idea as to the ease and sophistication of a system.

System Restrictions

The prospective Internet user should ascertain whether a provider has any applicable restrictions for whom service is available. For example, some providers might work only with not-for-profit organizations, or only with libraries located within a particular state, county, or city. There are also some technical aspects about which one should inquire. Chief among these are hours of availability, storage space limitations, and communications software requirements.

While in most cases none of these will prove to be a problem, it is prudent to check. Most providers make access available twenty-four hours per day, but some have regularly scheduled periods of downtime for maintenance or other purposes. Some providers impose upper limits on the amount of storage space available to a subscriber, or charge for storage space

beyond a certain level. The prospective user will generally find the upper limits to be adequate, or any additional charges beyond the limits to be acceptable. Most providers can handle any of the more popular communications packages and a lot of others as well; this will seldom be a problem, but it is worth confirming compatibility ahead of time.

Safeguards

In addition to understanding any limitations like the above, it is also important to know whether an Internet provider has taken any precautions against misuses of the system that could have negative repercussions for its users. Three of the major points in this regard include protection against computer viruses, general system security, and measures to ensure user confidentiality.

It may be difficult to protect all traffic into and out of a system against viruses, but the provider should at least run virus detection routines against any systems operations software it imports. On a broader scale, a prospective user may also want to inquire about the measures a provider takes to build security into the system and how often files are backed up. No systems that cater to external use like those of Internet providers will be totally secure from intrusion by unauthorized users, but the potential for harm can be minimized by a provider's having taken full advantage of security-minded measures that are available for most systems. And thirdly, the prospective user may want to question the provider about its policy on confidentiality of files and messages. The provider should have a clearly stated policy on under what circumstances any of a user's files or messages may be read by anyone other than the user or consulted by any outside party, without the express, situation-specific consent of the user.

User Support and Responsiveness

An area in which there is great variation among Internet providers is in the level and quality of the user support they offer. This can be a very important consideration in choosing among providers. There is nothing more frustrating in using a system than encountering a problem, being focused on finding the solution, and then running into a brick wall because there is no one available to help. Some Internet access providers offer little in the way of customer support other than a recorded message on an answering machine and an eventual call-back. On the other end of the spectrum, there are full-service providers who have a customer assistance hotline and

who offer proactive support in the form of users meetings, workshops, newsletters, and so forth.

In looking at a support structure, there are at least five basic areas that a prospective subscriber should look at:

- User documentation — Does the provider have a user manual for its system? Is it written in standard English or computerese? Does it adequately explain how to use all the functionalities (e-mail, telnet, ftp) that you are planning to use?
- Telephone support — Is there someone to call if you run into problems or have questions? What hours is telephone support available? Is there any additional charge for it?
- Training and workshops — Few providers' front-end interfaces are so simple that a new user cannot benefit from some initial training, and learning about resources available through the Internet is facilitated through educational opportunities. Does the provider offer start-up training? Advanced training? General and specialized workshops focusing on resources available on the Internet?
- Ongoing communication — Does the provider publish an ongoing newsletter or memos that discuss the use of the system? That include information about new resources on the Internet? Does the provider sponsor users meetings?
- Technical enhancements — Does the provider have a track record of making enhancements to its system that make it easier to use? When changes are made, how are users kept informed? Does the provider act on users' suggestions for improvements?

As with any complex electronic information system or service, a strong support structure can make an enormous qualitative difference. And Internet access through providers with high-quality support in place is not necessarily more expensive right now than access through providers that do not have strong support structures in place. This may change over time as providers are better able to assess their overall costs and more accurately price their services, but at least right now, users in most areas will find that one does not necessarily pay a high premium for quality support.

Conclusion

Prospective Internet users who find themselves in the position of having more than one readily available access option to choose from can compare the

alternatives from a variety of angles. Cost is certainly one important consideration, and there are a number of different factors that should be taken into account in understanding cost. The range of performance and support factors is equally broad. Even criteria like allegiance can play a major role.

For example, some users initially may tolerate less sophisticated features, particularly in a start-up mode, if the provider is a local consortium or company that in the long run may be more willing than others to cater to the interests of libraries and information centers or to a particular area of business concern.

Thus, there are many criteria for a prospective dial-up Internet user to consider in evaluating access options. The relative weighting of these will be a highly individual matter, but the most important thing is for prospective users to be aware of what to look for so that they can make a selection that will best meet their specific needs.

NOTES

1. Two articles last year that addressed specifically the use of the Internet by special librarians include Hope N. Tillman and Sharyn J. Ladner "Special Librarians and the Internet" *Special Libraries*, 83(Spring, 1992): 127-131; and Marian Bremer "Using the Internet in a 'Special' Library" *Internet World*, 3(October, 1992): 12-13. Those wanting more in-depth descriptions of

Internet capabilities will find a number of excellent guides available, among them: Ed Krol, *The Whole Internet: User's Guide & Catalog*, Sebastopol, Cal.: O'Reilly & Associates, 1992; Brendan P. Kehoe, *Zen and the Art of the Internet: A Beginner's Guide to the Internet*, 2nd ed., Englewood Cliffs, N.J.: Prentice Hall, 1993; Roy Tenant, John Ober, and Anne G. Lipow, *Crossing the Internet Threshold: An Instructional Handbook*, Berkeley, Cal.: Library Solutions Press, 1993; and Tracy L. LaQuey, *The Internet Companion: A Beginner's Guide to Global Networking*, Reading, Mass.: Addison-Wesley, 1992.

2. Notess, Greg. "Gaining Access to the Internet" *Online*, 17(September, 1992): 27-34.
3. One such list on the Internet is Peter Kaminski's Public Dialup Internet Access List, which may be obtained at kaminski@netcom.com with an e-mail message Send PDIAL. Lists available only on the Internet obviously are not of the greatest utility to those who do not yet have Internet access. Some of the sources listed in the first footnote above also contain lists of Internet access providers. Because of the pace at which new, especially local, providers are appearing on the scene, any list of this nature is far from inclusive.

Dennis Reynolds is president, CAPCON Library Network.

INTERNET NEWS

OCLC Minutes Available Via the Internet

Minutes from OCLC Users Council meetings, from October 1990 through the May 1993 meeting, are now available on the Internet. The **Users Council** comprises delegates from networks and service centers whose use of the OCLC System and contribution to the OCLC database qualify them for Users Council membership. **Delegates ratify amendments to the Articles of Incorporation and Code of Regulations of OCLC.**

The Users Council meets at OCLC three times a year to exchange ideas and to discuss issues important to librarians and educators using OCLC products and services.

To receive meeting minutes via the Internet, send a message to listserv@oclc.org. Commands should be typed on separate lines in the body of the message, not in the subject line.

Enter the command, index uc, to receive the index of archived minutes. Enter get [path] [file name] for the requested meeting minutes.

EXHIBIT 23

**IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE**

**INTELLECTUAL VENTURES I LLC and
INTELLECTUAL VENTURES II LLC,**

Plaintiffs,

v.

MOTOROLA MOBILITY LLC,

Defendant.

Case No. 1:11-cv-908-SLR

DECLARATION OF DR. RANDY HOWARD KATZ

1. My work address is 465 Soda Hall, RadLab, Computer Science Division, Electrical Engineering and Computer Science Department, University of California, Berkeley, California 94720-1776. Defendant Motorola Mobility LLC (“Motorola Mobility”) retained me in the above-referenced litigation. I am over the age of eighteen and am a citizen of the United States.

2. I am submitting this declaration, which sets out the opinions to which I would testify, if asked, in support of Motorola Mobility’s Motion for Summary Judgment of Invalidity with respect to U.S. Patent No. 7,409,450 (the “’450 Patent”).

3. I have extensive educational and professional experience in the fields of computer science and engineering, particularly with respect to networking technologies relevant to the ’450 Patent, which is more fully set out in paragraphs 2-10 of my expert report on the invalidity of the ’450 Patent, served on May 17, 2013

4. Based on my experience, by early 1998, there had been, and was ongoing, extensive research and development into networking systems, both wired and wireless, that prioritized data packets for transmission over a network. I was involved in some such efforts and have first-hand knowledge of systems that were published and/or actually used before early 1998, that prioritized packets, based on what they contained, for transmission over both wired and wireless networks. For example, I was on the technical advisory board for a company called Packeteer, which in 1996-97 developed a product known as Packet Shaper. Packet Shaper, which was first commercially introduced in the first quarter of 1997, Galloway Dep. Tr. 175:22-176:6, 101:24-102:10, classified data packets for transmission according to information contained within the data packets themselves, such as by analyzing the contents of headers of the data packets and/or the contents of the payload of the data packets. Katz Invalidity Rep., ¶¶ 160-163, 210-213. In addition, I am aware of and have reviewed papers and descriptions of systems developed by others that similarly prioritized packets for transmission over wireless networks.

5. U.S. Patent No. 6,463,096 (“Raleigh”), which was filed on June 12, 1998, describes one such system that prioritizes packets for transmission over shared wireless bandwidth from a CPE to a base station.

6. As described in my report at pages 23-115, 122-124, and 132-181, other systems that I have analyzed include the following:

- SWAN: A Mobile Multimedia Wireless Network by Prathima Agrawal, Eoin Hyden, Paul Krzyzanowski, Partho Mishra, Mani B. Srivastava, and John A. Trotter, April 1996, IEEE Personal Communications, Page(s): 18-33;
- The Packet Shaper system, which was developed and sold by Packeteer (for which I was a technical consultant on the technical advisory board) in the first quarter of 1997, as demonstrated in part by U.S. Provisional Patent Application No. 60/066,864 to Riddle et al (the “Riddle Provisional”);
- U.S. Patent No. 6,118,777 to Sylvain, which was filed on October 27, 1997;
- Japanese Patent Application Publication No. JP H08-107417, which was filed on October 5, 1994, and published on April 23, 1996;
- The Stelliga / Softcom system, which was described in U.S. Provisional Patent Application No. 60/090,939, filed on June 27, 1998, and U.S. Patent No. 6,625,650, filed on June 25, 1999;
- U.S. Patent No. 6,046,979 to Bauman, which was filed on May 4, 1998;
- U.S. Patent No. 6,005,851 to Craddock, which was filed on October 10, 1997;
- U.S. Patent No. 5,926,458 to Yin, which was filed on January 31, 1997; and
- U.S. Patent No. 6,125,397 to Yoshimura, which was filed on June 3, 1998.

7. While many different types of networking protocols abound, a number of protocols were in widespread use and understood by persons of ordinary skill in the art in early 1998. Among these well-known protocols were the Internet Protocol or “IP” protocol, the Transmission

Control Protocol or “TCP,” and the User Datagram Protocol or “UDP.” Each of these protocols – IP, TCP, and UDP – are examples of protocols of the type described as “packet-centric protocols” in the ’450 Patent. There were many types of networking equipment that used such protocols well before 1998, such as the IP router shown in Raleigh. I provided a more fulsome discussion of this component and the IP, TCP, and UDP protocols in paragraph 348 of my Invalidity Expert Report.

8. As I stated in my expert invalidity report, it is my opinion that Raleigh alone, or Raleigh in view of Packet Shaper, renders each of the asserted claims of the ’450 Patent invalid. Katz Invalidity Rep., p. 118-134

9. The following chart tracks the elements of the asserted claims and, based on my review of Dr. Gibson’s report disclosing his opinions, whether IV has disagreed that those elements are disclosed in the Raleigh, and/or Raleigh in view of Packet Shaper, prior art. The chart also provides illustrative citations to Raleigh and/or Packet Shaper, where each of the claim elements is disclosed.

Claim	Claim Element	Disclosing References	Disclosure: At Issue or Not At Issue.
1	A method comprising		
1a	coupling one or more subscriber customer premise equipment (CPE) stations with a base station over a shared wireless bandwidth using a packet-centric protocol	See, e.g.: Raleigh, 4:35-37, 4:40-49, 5:64-66, Fig. 1. Katz Invalidity Expert Report ¶¶ 346-349	Not At Issue. Dr. Gibson did not dispute that Raleigh discloses this claim element.

Claim	Claim Element	Disclosing References	Disclosure: At Issue or Not At Issue.
1b	allocating said wireless bandwidth and system resources based on contents of packets to be communicated over said wireless bandwidth, wherein the contents of each packet include a packet header and	See, e.g.: Raleigh 6:22-35, Fig. 3. Katz Invalidity Expert Report ¶¶ 350-351	Not At Issue. See Gibson Validity Report ¶ 176.
	wherein the allocating is responsive to at least one field in the packet header.	See, e.g.: Raleigh, 6:22-35, Fig. 3. Riddle Provisional, 5:29-6:2, 12:31-14:14, 15:19-17:25, Fig. 1D Katz Invalidity Expert Report ¶¶ 352-358	At Issue, as to motivation to combine Raleigh and Packet Shaper. See Gibson Validity Report ¶¶ 177-178.
2	The method of claim 1, wherein said packet-centric protocol comprises transmission control protocol/internet protocol.	See, e.g.: Raleigh 4:35-44, Fig. 3. Katz Invalidity Expert Report ¶ 361	Not At Issue. Dr. Gibson did not dispute that Raleigh discloses this claim element.
3	The method of claim 1, wherein said packet-centric protocol comprises user datagram protocol/internet protocol.	See, e.g.: Raleigh 4:35-44, Fig. 3. Katz Invalidity Expert Report ¶ 365	Not At Issue. Dr. Gibson did not dispute that Raleigh discloses this claim element.

Claim	Claim Element	Disclosing References	Disclosure: At Issue or Not At Issue.
5	The method of claim 1, wherein said coupling one or more subscriber CPE stations with said wireless base station comprises using a telecommunications access method including at least one of: a time division multiple access method; a time division multiple access/time division duplex access method; a code division multiple access method; and/or a frequency division multiple access method.	See, e.g.: Raleigh, 5:38-41, 6:38-41, Figs. 4A, 4B Katz Invalidity Expert Report ¶¶ 374-375	At Issue, as to motivation to combine Raleigh and Packet Shaper. See Gibson Validity Report ¶ 184.
8	The method of claim 1, wherein said allocating is further based on at least one of a packet header contents and a packet payload contents.	See, e.g.: Raleigh, 6:22-35, Fig. 3. Riddle Provisional, 15:26-16:5, 24:13-26:8. Katz Invalidity Expert Report ¶¶ 387-391	At Issue, as to motivation to combine Raleigh and Packet Shaper. See Gibson Validity Report ¶¶ 185-186
9	The method of claim 1 wherein the packets to be transmitted include packet-switched voice packets and data communication packets.	See, e.g.: Raleigh, 5:62-6:9, 6:30-31, Fig. 3. Katz Invalidity Expert Report ¶¶ 394-396	Not At Issue. Dr. Gibson did not dispute that Raleigh discloses this claim element.

10. Regarding the second portion of claim 1b, I noted in my report that the Raleigh system prioritizes packets by inspecting data in each packet to determine its priority. Dr. Gibson does not dispute my analysis on this point. Inspecting the contents of the packets necessarily

involves inspecting one or more of the different portions of the packet, which is made up of a header and a payload. My opinion remains that it would have been obvious to a person of skill in the art that prioritizing packets on the basis of what they contain entails looking at what they contain, which entails looking at either or both of the header or payload. To demonstrate that this technique was well known in the prior art, I cited to Packet Shaper as well as a number of other prior art systems that explicitly describe inspecting packet headers or packet payloads to prioritize packets for transmission, including some of those listed above in paragraph 6 of this declaration. Katz Invalidity Rep., ¶¶ 351-358, 387-391.

11. Inspection of the packets themselves is an analysis purely based on the specification for a particular packet format, such as the specification for IP, TCP, or UDP packets from their respective specifications, RFC 791, RFC 793, and RFC 768. See MMI-IV1048307-55 (RFC 791), MMI-IV1048356-444 (RFC 793), MMI-IV1048304-306 (RFC 768). These specifications dictate what information is stored where, in packet contents (e.g., what fields are included in the packet header), and are unrelated to the physical network being used. Both wired IP networks and wireless IP networks, for example, use IP packets of the same format. Thus, packet inspection is simply a basic networking issue that is resolved by examining the specification for the packets that are of interest to identify the appropriate header and/or payload fields to examine.

12. As I opined in my expert report, it would have been obvious to one of skill in the art to examine the contents of a packet header to determine a priority for the packet. Katz Invalidity Rep. ¶ 351. Packets contain headers and payload and inspecting the contents of a packet would have necessitated an inspection of one or more of these basic components of the packet. Further, because the contents of a header are well-defined by a specification, they provide an obvious starting place within the packet to inspect. In addition, as I noted in my report, examining packet contents to determine a corresponding priority was well known as were myriad techniques for analyzing packet contents, including packet headers and packet payloads. Katz Invalidity Rep., ¶¶ 351-358. Packet Shaper was one such system with which I was and remain familiar. In

addition, it provides illustrative prior art techniques for inspecting and analyzing the contents of packets to determine how to prioritize a packet for transmission. Packet Shaper is an excellent example of a system to which one of skill in the art would have turned for teachings related to analyzing packets to determine their priority.

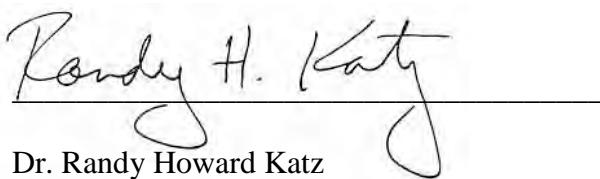
13. Regarding asserted claim 5, in my expert report I noted that Raleigh discloses different schemes for dividing bandwidth. As I previously noted, Katz Invalidity Rep., ¶¶ 372-375, Raleigh describes dividing upstream and downstream mediums by time and frequency. Dividing a wireless medium by time is referred to a time division multiple access (as recited in claim 5) and dividing a wireless medium by frequency is referred to as frequency division multiple access (as recited in claim 5). Thus, Raleigh does disclose at least these two access methods of claim 5, which requires only that one of the listed methods be used.

14. Regarding asserted claim 8, as I discussed above regarding the second portion of claim 1b, inspecting the contents of the packets will necessarily include inspecting one or more of the different portions of the packet, such as the header, the payload, or both. Thus, it likewise remains my opinion that given the very limited choices, examining either or both of the header or payload would have been necessary, or at the very least, obvious. To demonstrate that this technique was well known, I cited to Packet Shaper, which explicitly describes inspecting packet headers and packet payloads to prioritize packets for transmission. Katz Invalidity Rep. ¶¶ 388-391.

15. Lastly, as I stated in my expert report regarding secondary considerations of non-obviousness, I have reviewed Dr. Gibson's bases for secondary considerations of non-obviousness. His opinions that commercial success attributable to the patent may be found in the 3GPP and 3GPP2 standards lacks foundations. First, IV has dropped its allegations with respect to 3GPP, which undercuts its commercial success argument. Further, having reviewed Mr. Seely's expert report of non-infringement, I agree with his conclusion that compliance with the 3GPP2 standard does not necessitate infringement of the '450 Patent. Katz Secondary Considerations Report, ¶¶ 9-10. In fact, having reviewed Mr. Seely's analysis, I agree with his

conclusion that the accused Motorola products do not infringe the '450 Patent. Further, because Dr. Gibson has provided no relevant information regarding the licenses on which he relies as objective evidence of non-obviousness, I find that his opinion lacks any basis in evidence. Katz Secondary Considerations Report, ¶¶ 11-14. Thus, I maintain my opinion that Dr. Gibson has provided no objective evidence of any secondary considerations of non-obviousness.

Dated: August 19, 2013



Dr. Randy Howard Katz

Addendum to Declaration of Dr. Randy Howard Katz

Galloway deposition excerpts at Ex. 36, IA0790-0793.

Relevant Packet Shaper analysis at Ex. 37, IA0794-0805.

Representative excerpt of RFC791 at Ex. 38, IA0806-0810.

EXHIBIT 24

REDACTED
IN ITS
ENTIRETY

EXHIBIT 25

REDACTED
IN ITS
ENTIRETY

EXHIBIT 26

US006463096B1

(12) **United States Patent**
Raleigh et al.

(10) **Patent No.:** US 6,463,096 B1
(45) **Date of Patent:** Oct. 8, 2002

(54) **MAC PROTOCOL EMPLOYING MULTIPLE DATA RATES**

(75) Inventors: **Gregory G. Raleigh**, El Granada; **Vincent K. Jones**, Redwood Shores; **Michael A. Pollack**, Cupertino, all of CA (US)

(73) Assignee: **Cisco Systems, Inc**, San Jose, CA (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

(21) Appl. No.: **09/097,305**

(22) Filed: **Jun. 12, 1998**

(51) **Int. Cl. 7** **H04B 17/00**

(52) **U.S. Cl.** **375/225**; 375/221; 370/473

(58) **Field of Search** 375/140, 141, 375/257, 260, 295, 225, 263, 265; 370/329, 330, 327, 340, 341, 337, 347, 431, 436, 432, 464, 468, 471, 473, 475; 455/450, 464, 509, 451, 510

(56) **References Cited**

U.S. PATENT DOCUMENTS

5,481,542 A * 1/1996 Logston et al. 370/94.2
5,638,412 A * 6/1997 Blakeney, II et al. 375/377
5,751,708 A * 5/1998 Eng et al. 370/389
5,790,514 A * 8/1998 Marchok et al. 370/208
5,875,182 A * 2/1999 Hatzipapafotiou 370/321
6,031,832 A * 2/2000 Turina 370/348
6,069,883 A * 5/2000 Ejzak et al. 370/335
6,075,787 A * 6/2000 Bobeck et al. 370/395
6,088,335 A * 7/2000 I et al. 370/252
6,108,323 A * 8/2000 Gray 370/335
6,122,516 A * 9/2000 Thompson et al. 455/450
6,130,887 A * 10/2000 Dutta 370/347
6,157,623 A * 12/2000 Kerstein 370/315
6,167,270 A * 12/2000 Rezaifar et al. 455/442

OTHER PUBLICATIONS

N. Abramson, "The ALOHA System—Another Alternative for Computer Communications," 1970 Fall Joint Computer Conference, AFIPS Conference Proceedings, vol. 37.

Anthony S. Acampora, "An Introduction to Broadband Networks," 1994, Plenum Press, pp. 52-55.

Karol et al., "An Efficient Demand-Assignment Multiple Access Protocol for Wireless Packet Networks," 1995, ACM/Baltzer Wireless Networks, vol. 1, No. 3, pp. 267-179.

Jacobsen et al., "A Discrete Multitone-based Network Protocol for Multipoint-to-point Digital Communications in the CATV Reverse Channel," May 1995, Canadian Cable Television Association Technical Papers.

Östberg et al., "Multiaccess Methods based on OFDM," Telia Research AB.

Peled et al., "Frequency Domain Data Transmission Using Reduced Computational Complexity Algorithms," Apr. 1980, IEEE Int. Conference acoust., Speech, Signal Processing, Denver Co. pp. 964-967.

* cited by examiner

Primary Examiner—Chi Pham

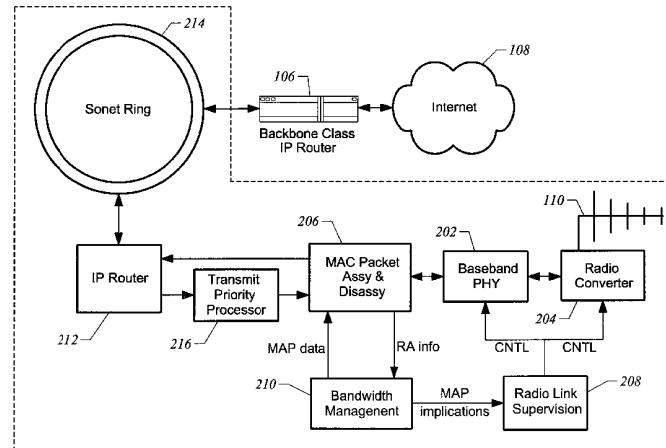
Assistant Examiner—Emmanuel Bayard

(74) **Attorney, Agent, or Firm**—Ritter, Lang & Kaplan LLP

(57) **ABSTRACT**

Method and apparatus for combining high data rate traffic and low data rate traffic on a common transmission medium while maximizing efficient use of available spectrum. Since spectrum is an economically valuable resource and transport of data generates revenue, the present invention directly leads to more profitable network operation. The disclosed systems are applicable to both wired and wireless transmission media. In one embodiment, a bandwidth reservation scheme provides that data rate may be varied so that when a particular data communication device is allocated a frame, it is also assigned a data rate for use in that frame. Because bandwidth usage varies with data rate, the division of available spectrum into channels for use by individual data communication devices may also vary among frames.

26 Claims, 7 Drawing Sheets



U.S. Patent

Oct. 8, 2002

Sheet 1 of 7

US 6,463,096 B1

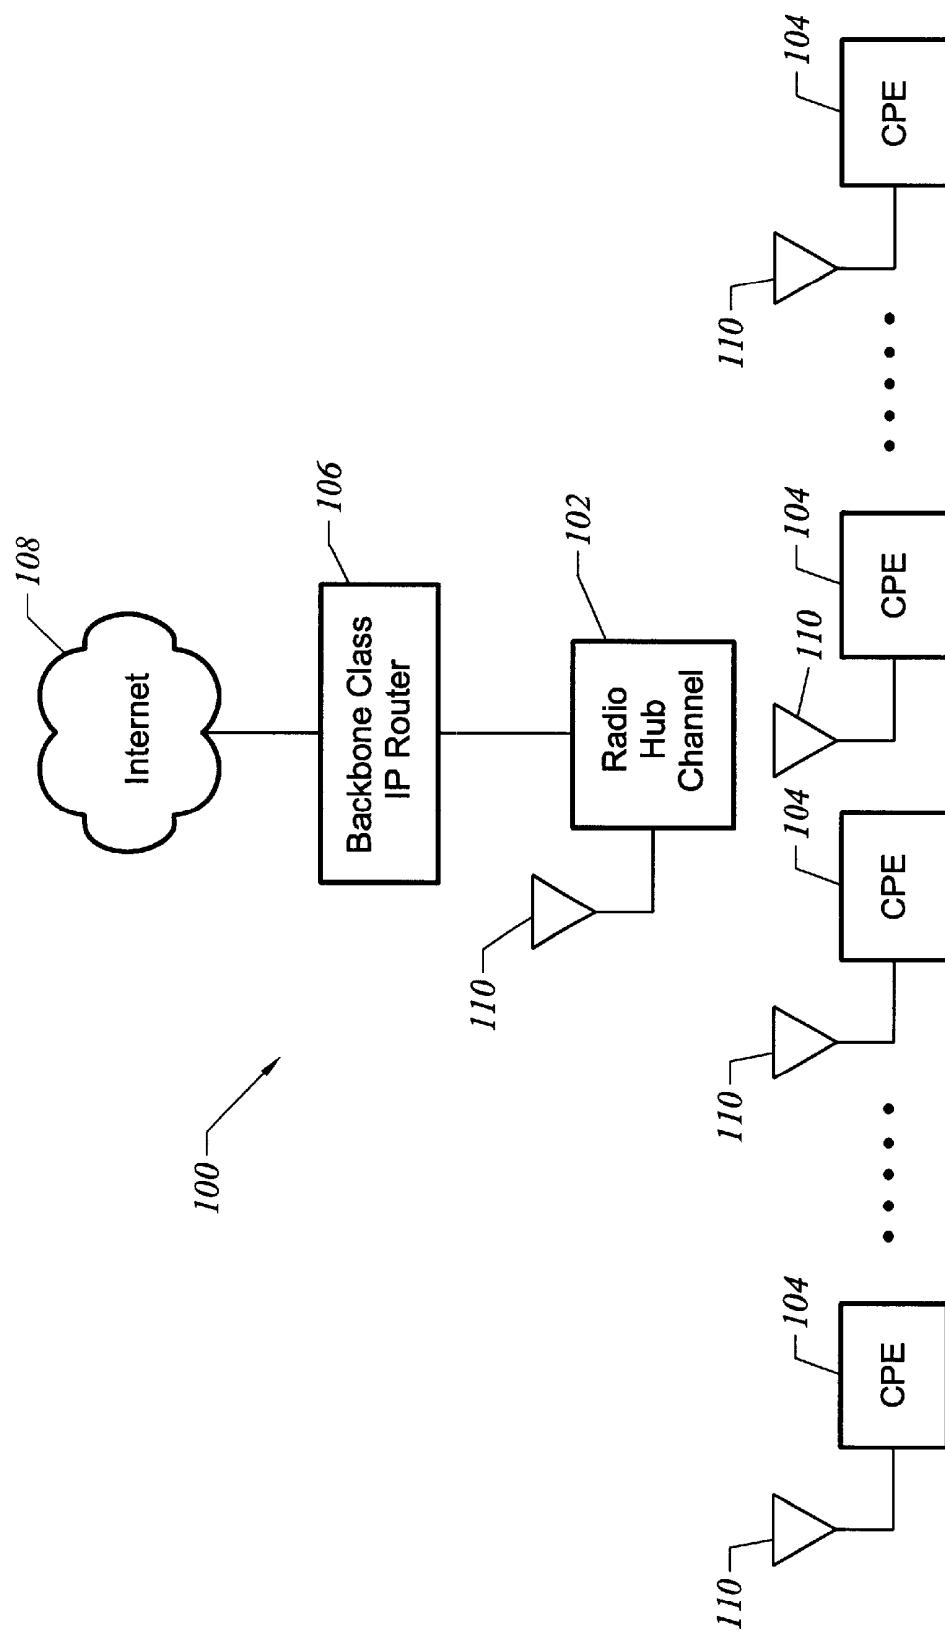


FIG. 1

U.S. Patent

Oct. 8, 2002

Sheet 2 of 7

US 6,463,096 B1

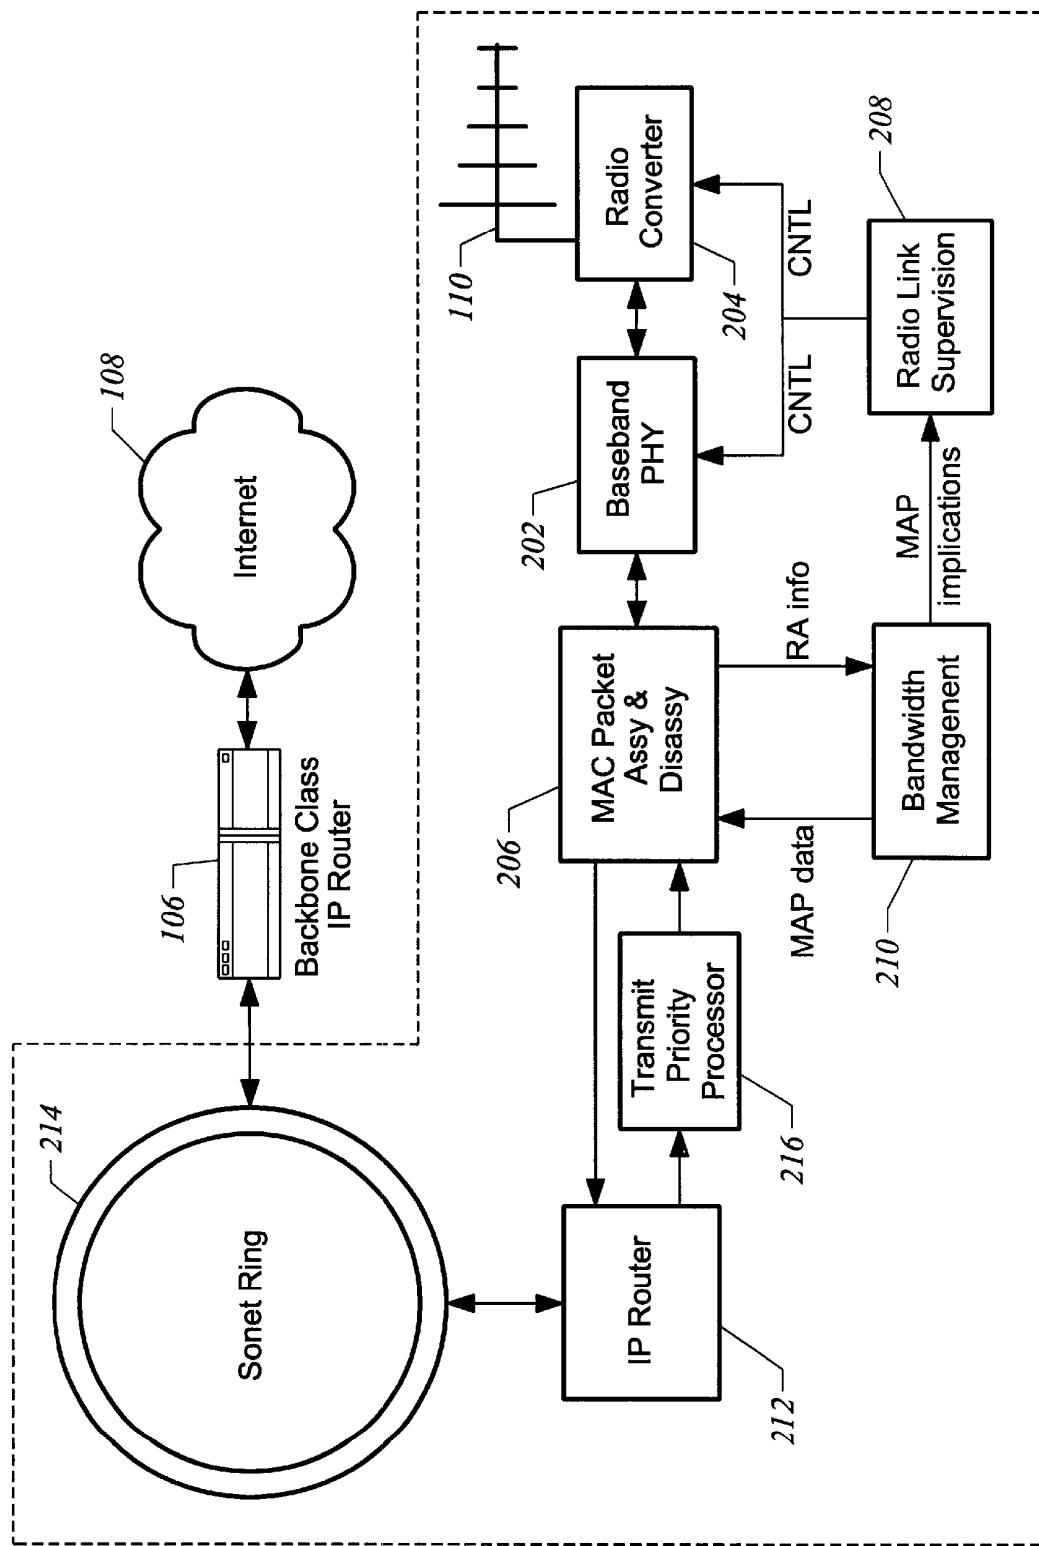


FIG. 2

U.S. Patent

Oct. 8, 2002

Sheet 3 of 7

US 6,463,096 B1

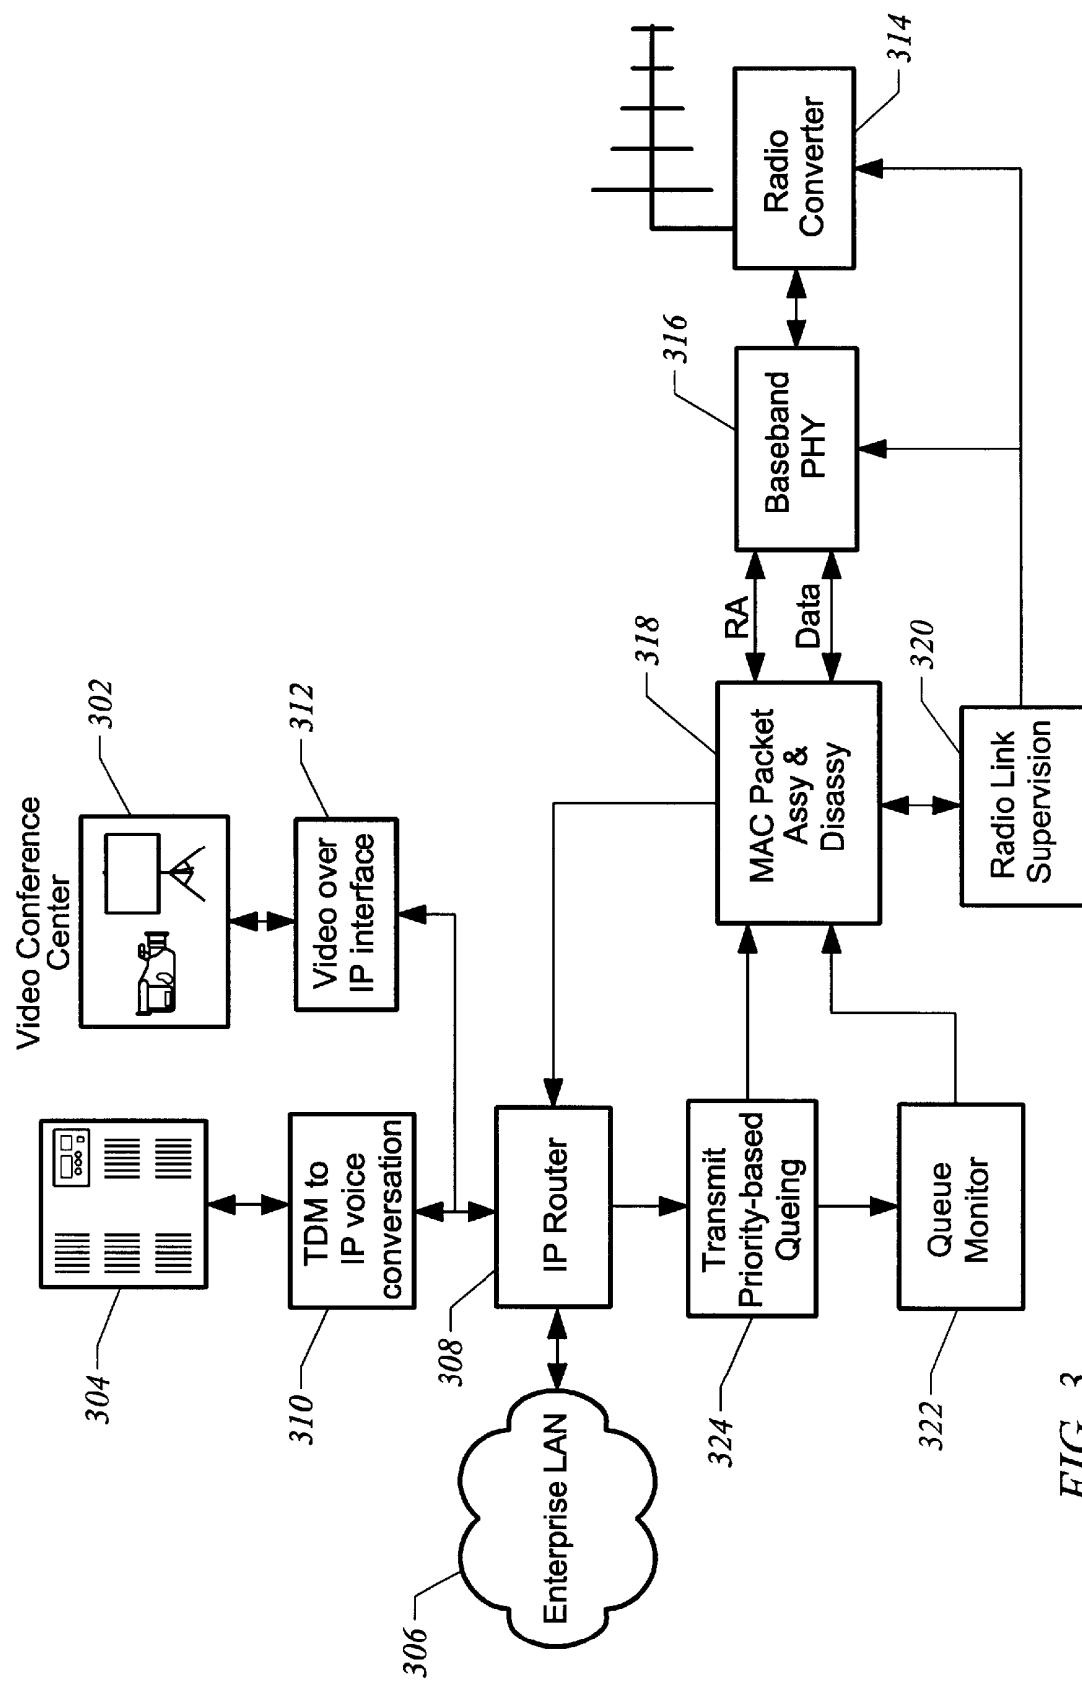


FIG. 3

U.S. Patent

Oct. 8, 2002

Sheet 4 of 7

US 6,463,096 B1

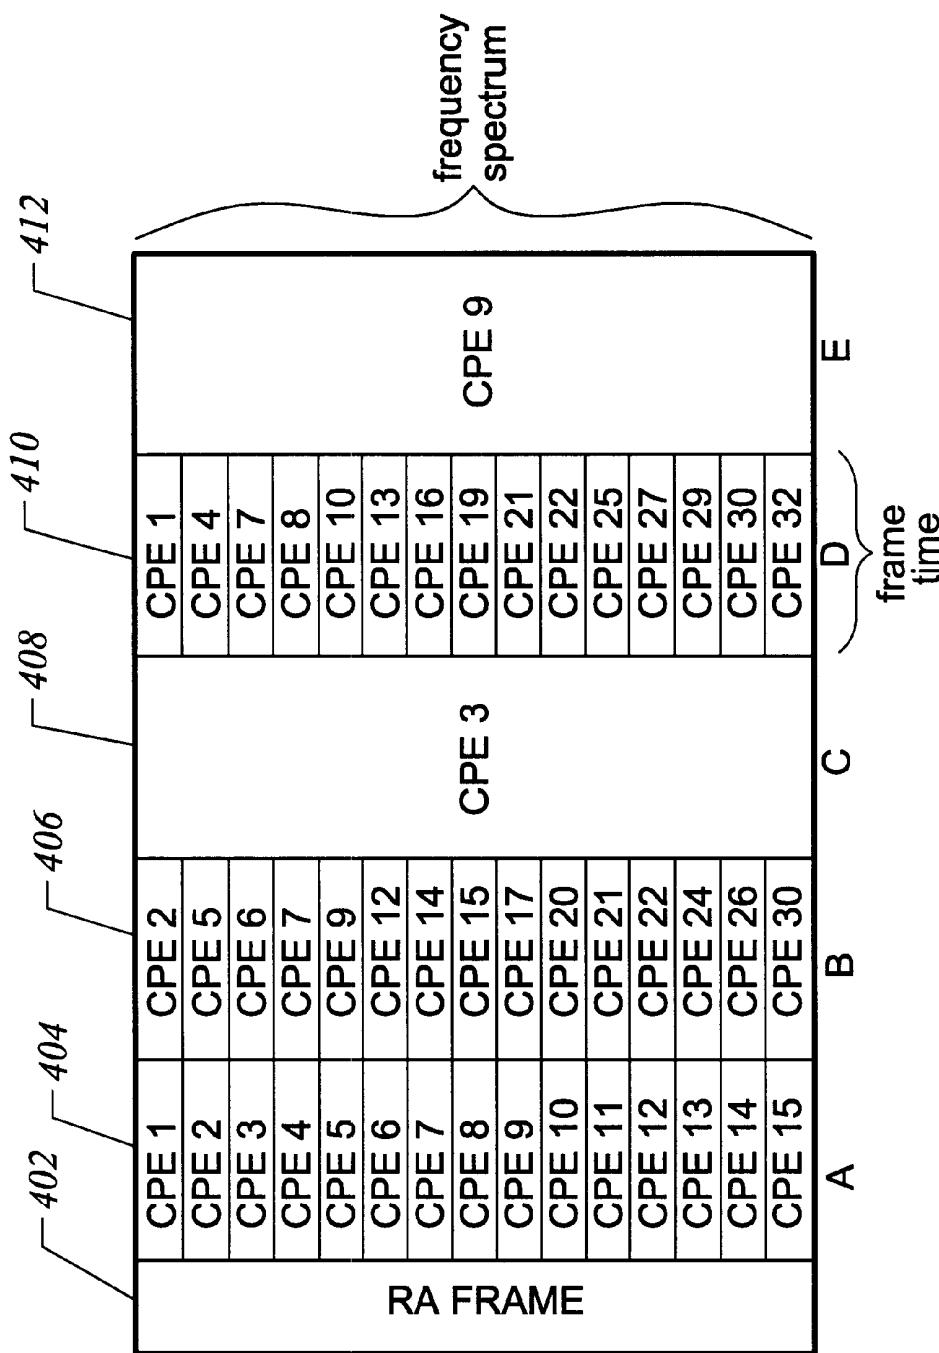


FIG. 4A

U.S. Patent

Oct. 8, 2002

Sheet 5 of 7

US 6,463,096 B1

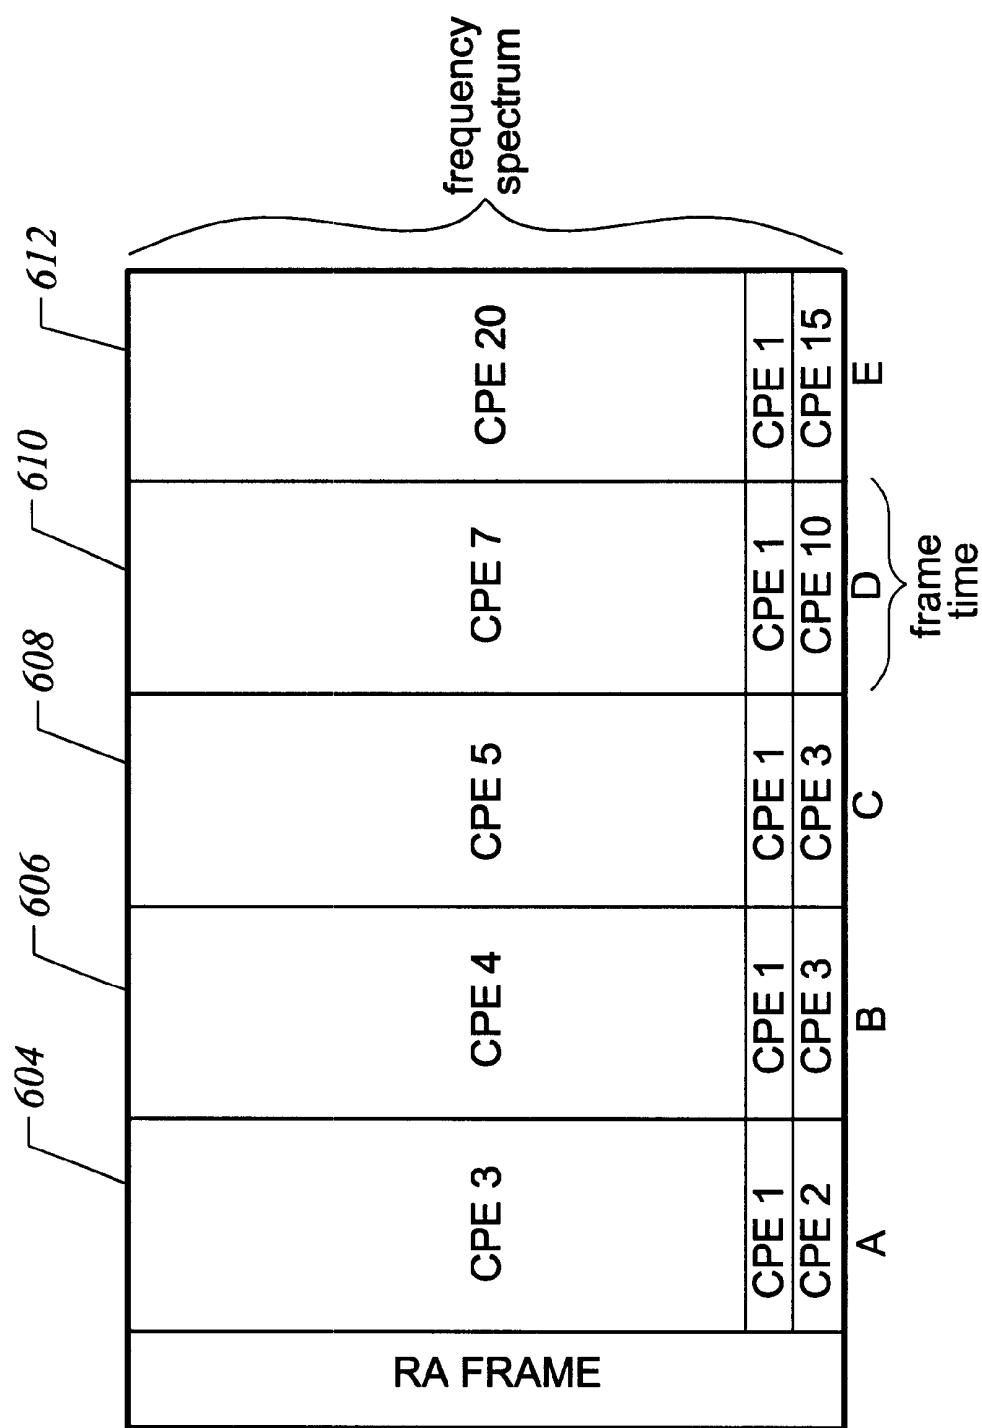


FIG. 4B

U.S. Patent

Oct. 8, 2002

Sheet 6 of 7

US 6,463,096 B1

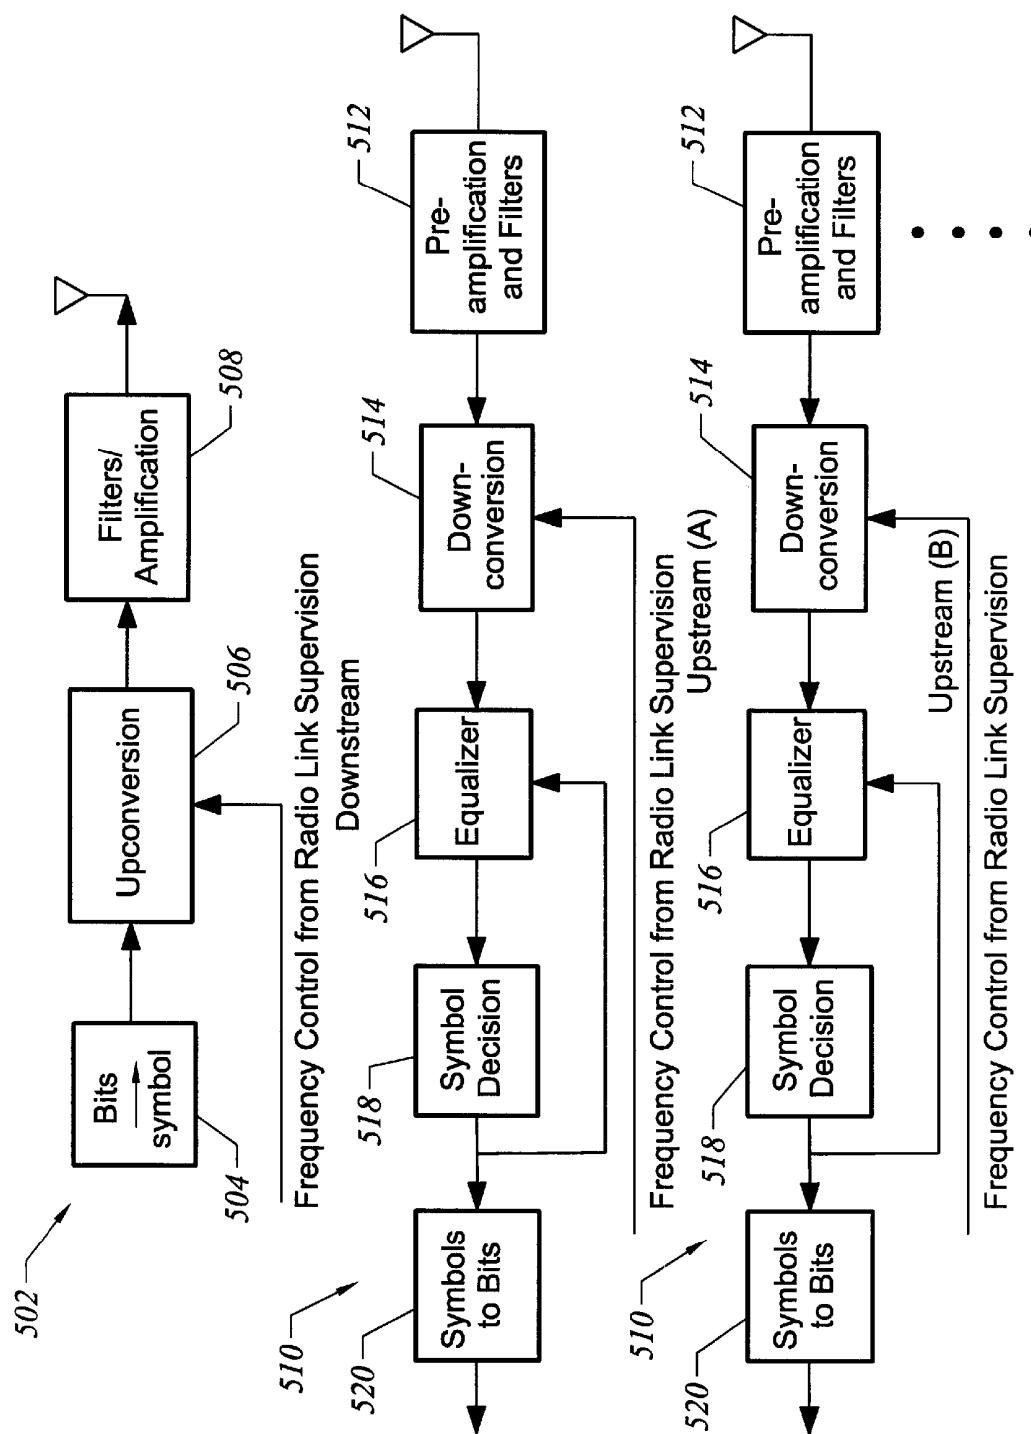


FIG. 5

U.S. Patent

Oct. 8, 2002

Sheet 7 of 7

US 6,463,096 B1

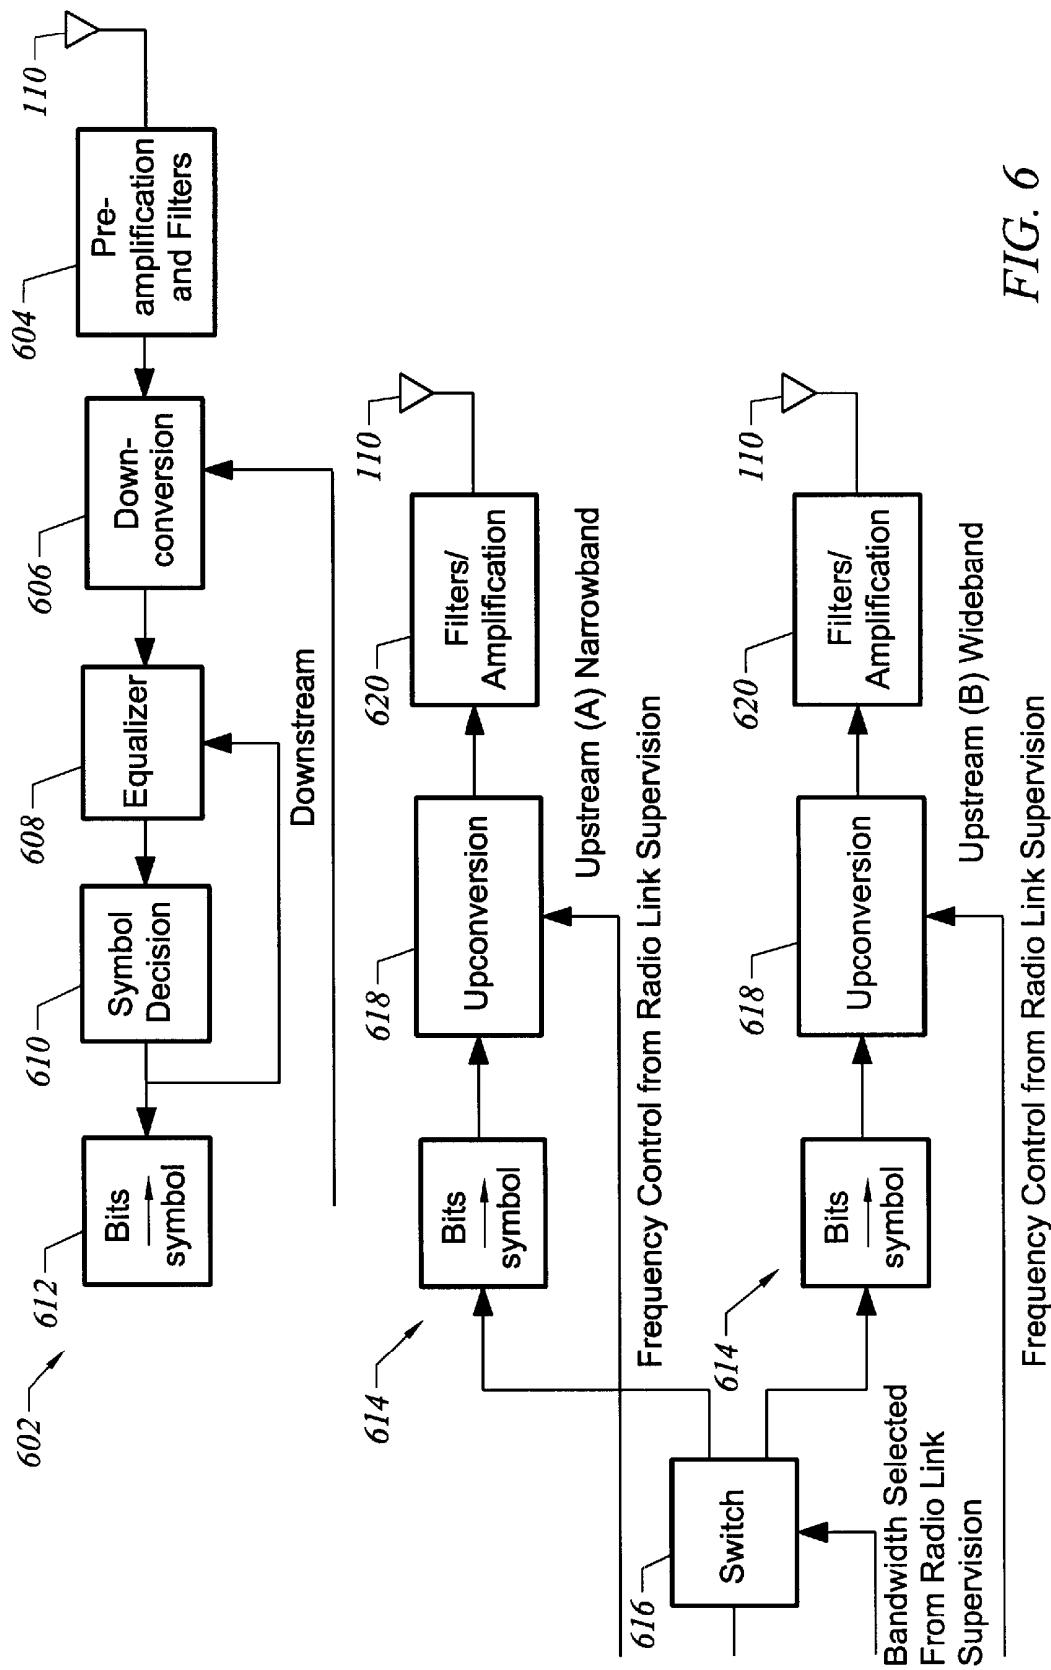


FIG. 6

US 6,463,096 B1

1

MAC PROTOCOL EMPLOYING MULTIPLE DATA RATES

BACKGROUND OF THE INVENTION

The present invention relates to digital communication and more particularly to protocols for communicating data originating from sources having disparate data rates over a shared medium.

Trends in digital communication point toward a common transmission medium providing both high data rate services such as digital video and low data rate services such as voice. Internet access is inherently a mixed service. Upstream requests for information typically include minimal data while downstream traffic may include graphics or even live video.

Specific examples of such a common transmission medium include a wireless local loop (WLL) that substitutes for the local telephone loop and provides additional high data rate services such as video and Internet access. Another example is a CATV network that has been updated to provide high data rate services and voice service.

A key objective is maximizing efficiency in use of bandwidth. The available bandwidth is shared among multiple data communication devices. When a data communication device is allocated all or part of the available bandwidth, it should make efficient use of its allocation. Depending on the protocols and modulation systems used, a certain percentage of the data is devoted to network operation rather than customer service. This is referred to as overhead. Consider a network where packets of information are communicated in successive frames and:

d_p =Payload data (bits)—the number of payload bits contained in a frame,

r =Data rate (bits/sec), proportional to spectrum used (Hz).

Data rate refers to the rate at which information is communicated through the wireless medium. Information rate roughly represents the rate of generation of payload data.

t_f =Frame time (sec)—the duration of the smallest unit of time that may be allocated to a data communication device for transmission on the shared medium. Note that a packet of like data, such as voice or data, may be transmitted in a single frame or may be divided among many frames.

t_g =Overhead time (sec), including guard time, training, and synchronization that is required for each frame.

The system efficiency associated with the overhead given by t_g is

$$eff = \frac{d_p}{r(t_g + t_f)}.$$

This value of efficiency reaches 100% when the overhead time is zero and the frame time equals the payload data divided by the data rate (when the frame time is exactly the time required to transmit the payload data at the transmission rate).

The network designer is left free to vary frame time to maximize efficiency. However, it is difficult to reconcile the needs of different traffic types. Consider choosing one frame time to accommodate both low information rate voice traffic and high information rate data traffic. Due to transmission latency requirements, voice traffic requires frequent frame transmissions to reduce latency. Hence, voice requires a

2

short frame time. Furthermore, the amount of data to be sent in these frames is small since voice is low information rate. If long frames are sent, voice traffic is insufficient to fill each frame, resulting in wasted bandwidth. However, sending small frames incurs a different type of bandwidth loss. The fixed overhead associated with each frame substantially reduces spectral efficiency. This becomes particularly significant for high information rate traffic, where data must be divided over many frames instead of being efficiently transmitted in long frames. This conflict can be described in an example.

Consider choosing a frame time to efficiently transmit both a 64-byte voice packet and a 1000-byte data packet. Assume an overhead time of 3 us ($t_g = 3$ us), a data rate of 30 Mbits/sec, and two candidate frame times of 17 us and 267 us. For a frame time of 17 us, the efficiency for a 64-byte packet of data is,

$$eff_{64} = \frac{d_p}{r(t_g + t_f)} = \frac{512}{30 \times 10^6 ((3 + 17) \times 10^{-6})} = \frac{17}{3 + 17} = 85\%.$$

Since the frame time is exactly the amount of time required to transmit 64 bytes at a 30 Mbit/sec rate, this is the maximum efficiency at this data rate. A 1000-byte packet would be spread among 64-byte transmission opportunities corresponding to individual frames. Hence the efficiency for a 1000 byte packet is approximately the same as for the 64-byte packet. (To be precise, the efficiency is slightly less than 85% since the final frame is not fully utilized.)

Increasing the frame time can increase this efficiency by reducing the overhead. For example, a frame time of 267 us results in close to 99% efficiency for a 1000 byte packet,

$$eff_{1000} = \frac{d_p}{r(t_g + t_f)} = \frac{8000}{30 \times 10^6 ((3 + 267) \times 10^{-6})} = \frac{267}{3 + 267} = 99\%.$$

Unfortunately, this large frame time causes severe inefficiency for the 64-byte packet because a large portion of the frame is left unutilized. Here it is assumed that because of latency requirements, it is not feasible to collect multiple 64-byte packets to fill a frame. For example, it would take a 64 kbps voice source over 125 ms to fill a 1000-byte frame, which results in intolerable latency. Allowing 8 ms of latency for the collection of one 64-byte packet, the long frame capable of supporting 1000 bytes which carries one 64-byte packet has very poor efficiency:

$$eff_{64} = \frac{d_p}{r(t_g + t_f)} = \frac{512}{30 \times 10^6 ((3 + 267) \times 10^{-6})} = \frac{17}{3 + 267} = 6\%.$$

No single choice of frame time leads to efficient use of the spectrum for both high data rate and low data rate traffic.

SUMMARY OF THE INVENTION

The present invention provides methods and apparatus for combining high data rate traffic and low data rate traffic on a common transmission medium while maximizing efficient use of available spectrum. Since spectrum is an economically valuable resource and transport of data generates revenue, the present invention directly leads to more profitable network operation. The systems and methods provided by the present invention are applicable to both wired and wireless transmission media. In one embodiment, a bandwidth reservation scheme provides that data rate may be varied so that when a particular data communication

US 6,463,096 B1

3

device is allocated a frame, it is also assigned a data rate for use in that frame. Because bandwidth usage scales with data rate, individual data communication devices will be assigned to possibly different spectrum bandwidth on a frame-by-frame basis.

A first aspect of the present invention provides a method for allocating access to a common transmission medium among a plurality of data communication devices. The method includes steps of: assigning a transmission frame to a particular data communication device, assigning a data rate for the particular data communication device to employ in the transmission frame, and transmitting the transmission frame assignment and the data rate assignment to the particular data communication device.

A second aspect of the present invention provides an alternative method for allocating access to a common transmission medium among a plurality of data communication devices. The method includes steps of: receiving access request messages from requesting data communication devices at a hub, the access request messages requesting access to the common transmission medium, in response to the access request messages, at the hub, allocating access to the common transmission medium in both the frequency and time domain among the requesting data communication devices, and thereafter transmitting from the hub to the requesting access devices, instructions for each requesting access device to transmit at particular times, and at particular data rates chosen according to the allocating step.

A third aspect of the present invention provides a digital communication network including a plurality of data communications devices transmitting via a common transmission medium, and a hub receiving signals from the data communications devices via the common transmission medium. The hub includes: a bandwidth manager that receives access request messages from requesting data communication devices, the access request messages requesting access to the common transmission medium, and that allocates access to the common transmission medium in both the frequency and time domain among the requesting data communication devices. The hub further includes a link supervisor that transmits from the hub to the requesting access devices, instructions for each requesting access device to transmit at particular, and at particular data rates chosen in accordance with allocations by the bandwidth manager.

A fourth aspect of the present invention provides a hub. The hub includes: a receiver that receives signals from the data communications devices via the common transmission medium, a bandwidth manager that receives access request messages from requesting data communication devices, the access request messages requesting access to the common transmission medium, and that allocates access to the common transmission medium in both the frequency and time domain among the requesting data communication devices. The hub further includes a link supervisor that transmits from the hub to the requesting access devices, instructions for each requesting access device to transmit at particular times, and at particular data rates chosen in accordance with allocations by the bandwidth manager.

A fifth aspect of the present invention provides a data communications device for use in a network. The data communications device includes a bandwidth manager that transmits requests for access to a common transmission medium to a hub. The data communications device further includes a link supervisor that receives medium access instructions from the hub, the medium access instructions

4

specifying data rate, transmission time, and transmission frequency for transmissions to the data communications device, and that controls transmission of information via the common transmission medium in accordance with the specified data rate, transmission time, and transmission frequency.

A further understanding of the nature and advantages of the inventions herein may be realized by reference to the remaining portions of the specification and the attached drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 depicts a representative communications network suitable for implementing one embodiment of the present invention.

FIG. 2 depicts a hub for the network of FIG. 1 according to one embodiment of the present invention.

FIG. 3 depicts a data communications device in the network of FIG. 1 according to one embodiment of the present invention.

FIG. 4A depicts frame assignments in the time domain according to one embodiment of the present invention.

FIG. 4B depicts channel assignments in the frequency domain in a representative frame according to one embodiment of the present invention.

FIG. 5 depicts a hub radio system according to one embodiment of the present invention.

FIG. 6 depicts a data communications device radio system according to one embodiment of the present invention.

DESCRIPTION OF SPECIFIC EMBODIMENTS

FIG. 1 depicts a representative communications network 100 suitable for implementing one embodiment of the present invention. Although communications network 100 is depicted as a wireless network, it will be understood that the present invention is applicable to both wired and wireless networks. A hub 102 acts as a central access point for network 100. Hub 102 may communicate with a plurality of CPEs 104 (customer premise equipment) which represent data communication devices. A backbone class IP router 106 may interconnect hub 102 and the Internet 108.

In one embodiment, network 100 implements a wireless local loop that provides local telephone services as well as high data rate services. Antennas 110 are used to couple hub 102 and CPEs 104 to the common wireless transmission medium. Spectrum is allocated for the use of network 100. In a preferred embodiment, downstream communications (i.e., communications from hub 102 to CPEs 104) are frequency domain duplexed with upstream communications (from CPEs 104 to hub 102). However, downstream and upstream traffic may also be duplexed in the time domain.

In another embodiment, network 100 is a cable television system. Hub 102 represents the cable head-end and CPEs 104 represent subscriber units coupled to a common coaxial transmission medium interconnecting the networks. The present invention does not assume a particular modulation system. Representative modulation systems include QAM and OFDM.

FIG. 2 depicts further details of hub 102. In order to interact with the common transmission medium, hub 102 includes a baseband physical layer processor 202 and in some applications such as wireless or cable, a radio converter 204. Baseband physical layer processor 202 includes hardware for implementing error correction coding, and any

US 6,463,096 B1

5

particular modulation scheme employed such as OFDM or QAM. It is understood that baseband physical layer processor 202 includes hardware and software to support multiple upstream receivers and at least one downstream transmitter. For transmission, radio converter 204 converts the baseband output of baseband physical layer processor 202 to a selected frequency for transmission. For reception, radio converter 204 converts a signal received at a selected reception frequency to baseband for input to baseband physical layer processor 202. Again, it is understood that radio converter 204 includes hardware and software to support multiple upstream receivers and at least one downstream transmitter. Further exemplary detail of these two stages is discussed with reference to FIG. 5.

A MAC processor 206 is responsible for multi-access control processing. MAC processor 206 receives and transmits IP packets to other components of hub 102. MAC processor 206 packages the IP data from the transmit priority processor 216 within the packet format specified by the operative MAC (medium access control) protocol. MAC processor 206 extracts IP data from MAC packets received from baseband physical layer processor 202 and forms the IP packets transmitted on the other end of the link (CPEs). Some of the MAC packets received by MAC processor 206 include access requests from CPEs 104. These access requests are forwarded to a bandwidth management processor 210.

In response to the access requests, a bandwidth management processor 210 allocates available upstream bandwidth among CPEs 104. Any scheduling technique may be used according to the present invention. One consideration in assigning center-frequency is channel quality available at different center-frequencies taking into account signal to noise ratio and/or signal to noise-plus-interference ratio. The bandwidth management processor 210 forwards assignments of frequency, data rate, and transmission frame to MAC processor 206 for inclusion in MAC packets to be transmitted downstream. These assignments of center-frequency, data rate, and transmission frame provide the CPEs the information regarding the time-frequency division of the upstream medium. This information is referred to as the MAP.

This MAP information is also forwarded to the radio link supervision processor 208. The radio link supervisor 208 partitions the baseband physical layer processor 202 and the radio converter 204 for proper reception of the upstream according to the MAP information. The operation of bandwidth management processor 210 and radio link supervision processor 208 is highly related. They may operate as either independent or integrated software packages on the same computer system.

An IP router 212 exchanges IP packets with backbone class IP router 106 via a Sonet ring 214. IP router 212 receives the IP packets from the MAC processor 206. IP packets to be transmitted out over network 100 are prioritized by a transmit priority processor 216. For example, voice packets and other real-time data are given a higher priority than other kinds of data. Priority processor 216 queues up IP packets to be transmitted and forwards them to MAC processor 206 in order of priority.

FIG. 3 depicts a representative CPE 104 as an example of a data communications device in the network of FIG. 1 according to one embodiment of the present invention. A video conference center 302, a PBX 304, and an Enterprise LAN 306 are representative data sources and destinations. An IP router 308 is connected to Enterprise LAN 306, to

6

PBX 304 via a TDM to IP interface 310, and to video conference interface 302 via a video over IP interface 312. A radio converter 314 and a baseband physical layer processor 316 essentially mirror the similarly named components of hub 102. A radio converter 314 and a baseband physical layer processor 316 include hardware and software to support at least one downstream receiver and at least two upstream transmitters or a single upstream transmitter capable of varying its data rate.

A MAC processor 318 operates to assemble and dissemble packets conforming to the operant MAC protocol. Much of the data extracted from the received MAC packets is in the form of IP packets which are forwarded to IP router 308. Some of the extracted data includes the MAP which carries instructions assigning transmission center-frequencies, data rates and frame times. These instructions are forwarded to a radio link supervision processor 320. Radio link supervision processor 320 controls the data rate, transmission times and center-frequencies of operation for baseband physical layer processor 316 and radio converter 314.

A queue monitor 322 originates requests for access to the common transmission medium. These access requests are forwarded to MAC processor 318 for forwarding to hub 102. The requests include the amount and priority of information to be transmitted. IP packets to be transmitted are forwarded to MAC processor 318 from transmit priority processor 324. Transmit priority processor 324 receives packets from IP router 308 that are to be directed to hub 102 and queues them in order of priority. Again, voice and other real-time traffic is given higher priority. Transmit priority processor 324 also indicates when data is to be transmitted and the amount of data to be transmitted to queue monitor 322. It is in response to these inputs that queue monitor 322 generates access requests.

FIG. 4A depicts a MAP with frame, data rate and center-frequency assignments according to one embodiment of the present invention. According to the present invention, the spectrum available for downstream communications is divisible in both the frequency and time domains. FIG. 4A shows a series of frames in the time domain. A frame is here understood to be a unit of time for which access to the common transmission medium may be assigned to one or more CPEs 104. A request access (RA) frame 402 is where individual CPEs may request access to the common transmission medium. Any known MAC scheme may be used to control access to the medium in this frame such as CSMA, CSMA/CD, etc. If RA frame 402 includes an OFDM burst, access contention during RA frame 402 may be ameliorated by assigning different groups of OFDM tones for use by different CPEs 104. This technique is explained in greater detail in the U.S. patent application Ser. No. 09/019,938 entitled MEDIUM ACCESS CONTROL PROTOCOL FOR OFDM WIRELESS NETWORKS filed on Feb. 6, 1998. This application is assigned to the assignees of the present application and its contents are herein incorporated by reference.

In the illustrated example, each CPE 104 may transmit upstream during a given frame at either a 2 Mbps data rate, a 30 Mbps data rate, or not at all. The present invention is, however, not restricted to any particular data rate, or number of possible data rates, or mixtures of data rates in a frame. In an A frame 404, 15 CPEs 1-15 are scheduled to transmit, each transmitting at 2 Mbps at differing center-frequencies. In a B frame 406, a different set of CPEs 2, 5, 6, 7, 9, 12, 14, 15, 17, 20, 21, 22, 24, 26, and 30 are scheduled to transmit. In a C frame 408, the entire upstream spectrum is

US 6,463,096 B1

7

reserved for a single CPE 3 which transmits at 30 Mbps. In a D frame 410, the upstream spectrum is again divided among 15 CPEs 1, 4, 7, 8, 10, 13, 16, 19, 21, 22, 25, 27, 29, 30, and 32. In an E frame 412, a single CPE 9 occupies the entire upstream spectrum. Thus, many CPEs may simultaneously transmit as low data rate sources or one CPE may transmit at a high data rate. The fluctuations of network traffic are typically such that demand for high data rate transmission is episodic meaning that intermittent high information rate sources may be serviced adequately without compromising the needs of low information rate sources. Assigning frames such as in FIG. 4A is done by bandwidth management processor 210.

In a preferred embodiment, the duration of the 30 Mbps frames is 267 us and each such frame holds a 1000 byte MAC layer packet. The duration of the 2 Mbps frames is 256 us and each CPE transmitting in such a frame transmits a 64 byte MAC layer packet. In an alternative embodiment, frame length is the same for both data rates. Preferably, downstream communications do not share the spectrum employed for upstream communications. However, the downstream communication may be multiplexed in the time domain with the upstream communication. Frames or frequency slots within frames may then be allocated to downstream transmission.

In another preferred embodiment, each CPE 104 may transmit upstream during a given frame at either a 2 Mbps data rate or a 26 Mbps data rate, or not at all. However, the MAP assigns the same number of distinct data rate slots for every frame. This is shown in FIG. 4B, where each frame consists of two 2 Mbps frequency slots and one 26 Mbps frequency slot. This MAP construction simplifies the baseband physical layer processor 202 and 316, along with the radio converter 204 and 314. In an A frame 604, 2 CPEs 1-2 are scheduled to transmit, each transmitting at 2 Mbps at differing center-frequencies. In this same A frame, CPE 3 is scheduled to transmit at 26 Mbps data rate. In a B frame 406, CPEs 1 and 3 are scheduled to transmit at 2 Mbps while CPE 4 transmits at 26 Mbps. Frames C, D and E are used by other CPEs. Note that CPE 1, by use of the MAP, has been allocated a constant data rate channel of 2 Mbps. This MAP assignment is done by bandwidth management processor 210.

FIG. 5 depicts details of baseband physical processor 202 and radio converter 204 of hub 102. These details are presented for a QAM application, although the present invention is not limited to any particular modulation system. A single downstream system 502 is depicted. Bits corresponding to MAC packet contents are received by bits to symbol converter 504 and mapped to appropriate positions on a QAM constellation. An upconversion stage 506 converts the signal output of bits-to-symbol converter 504 to the frequency allocated for downstream transmission. This frequency may be selected by a signal from radio link supervision processor 208. An amplification and filter stage 508 then outputs the signal onto the common transmission medium. There may of course be more than one such downstream system.

A series of upstream receiver systems 510 are provided. In one embodiment, there are upstream receivers for each possible 2 Mbps channel and a separate receiver for use in frames in which the entire upstream spectrum includes a single 30 Mbps channel. Alternatively, there may be a series of upstream receiver systems 510 with selectable or fully variable data rate. In another embodiment, the entire upstream spectrum is digitized and processed appropriately for each frame.

8

Each depicted upstream receiver system includes a preamplification and filter stage 512 for receiving a signal from the common transmission medium. A downconversion stage 514 converts the received signal to baseband. Downconversion stage 514 sets the frequency to be received based on frequency control input from radio link supervision processor 208. An equalizer 516 seeks to correct for distortion in the transmission medium. A symbol decision stage 518 estimates the transmitted symbols based on the output of equalizer 516. Equalizer 516 is preferably a decision feedback equalizer (DFE) and adapts in response to the output of symbol decision stage 518. A symbol to bits conversion stage 520 then generates the contents of the MAC packets.

FIG. 6 depicts representative details of baseband physical layer processor 316 and radio converter 314 of CPEs 104. Again, the details are presented in the context of QAM. A downstream receiver system 602 is used to receive data from hub 102. A preamplification and filter stage 604 extracts the signal from the common transmission medium. A downconversion stage 606 converts the received signal to baseband. The reception frequency may be controlled by a signal from radio link supervision processor 320. An equalizer 608 corrects for channel distortion in the transmission medium. A symbol decision stage 610 recovers the transmitted symbols. Equalizer 608 is preferably a DFE that adapts in response to the output of symbol decision stage 610. A symbol to bit conversion stage 612 then converts the received symbols to the MAC packet data.

Two upstream transmitter systems 614 are depicted, one for the 2 Mbps transmission data rate and one for 30 Mbps. Components may be shared between the two systems. In an alternative embodiment, a single transmitter system with selectable or even fully variable data rate is used.

In the depicted system, a switch 616 selects which bandwidth is to be used during a particular frame. MAC packet data is sent to the appropriate upstream transmitter system based on the setting of switch 616. A bit to symbol conversion stage 618 converts the MAC packet data to QAM symbols. An upconversion stage 620 sets the transmission frequency based on input from radio link supervision processor 320. An amplification and filter stage 622 prepares the signal for transmission via the common transmission medium.

The multirate upstream model provided by the present invention provides high quality service to both low information rate and high information rate sources while maximizing spectral efficiency. Also, latency is minimized for real-time traffic such as voice and video conference data.

It is understood that the examples and embodiments described herein are for illustrative purposes only and that various modifications or changes in light thereof will be suggested to persons skilled in the art and are to be included within the spirit and purview of this application and scope of the appended claims. All publications, patents, and patent applications cited herein are hereby incorporated by reference.

What is claimed is:

1. In a data communication network, a method for allocating access to a common transmission medium among a plurality of data communication devices comprising:

assigning a transmission frame to a particular data communication device wherein others of said data communication devices do not transmit in said transmission frame;

assigning a data rate for said particular data communication device to employ in said transmission frame; and

US 6,463,096 B1

9

transmitting said transmission frame assignment and said data rate assignment to said particular data communication device.

2. The method of claim 1 further comprising: transmitting an information frame from said data communication device during said transmission frame and at said data rate.

3. The method of claim 1 wherein assigning a data rate comprises:

selecting said data rate from two available data rates.

4. The method of claim 1 wherein assigning a data rate comprises:

selecting said data rate from among at least three available data rates.

5. The method of claim 1 further comprising: transmitting an access request from said particular data communication device to a hub, said access request indicating an amount of data to be transmitted by said particular data communication device.

6. The method of claim 1 wherein transmitting comprises 20 transmitting a center-frequency assignment to said particular data communication device.

7. The method of claim 6 wherein said center-frequency assignment is selected based on available transmission quality at a plurality of transmission frequencies.

8. In a data communication network, a method for allocating access to a common transmission medium among a plurality of data communication devices, said method comprising:

receiving access request messages from requesting data communication devices at a hub, said access request messages requesting access to said common transmission medium;

in response to said access request messages, at said hub, allocating access to said common transmission medium in both the frequency and time domain among said requesting data communication devices; and thereafter transmitting from said hub to said requesting access devices, instructions for each requesting access device to transmit at particular times and frequencies, and at particular data rates chosen according to said allocating step.

9. A digital communications network comprising: a plurality of data communications devices transmitting via a common transmission medium; and

45 a hub receiving signals from said data communications devices via said common transmission medium, said hub comprising:

a bandwidth manager that receives access request messages from requesting data communication devices, said access request messages requesting access to said common transmission medium, and that allocates access to said common transmission medium in both the frequency and time domain among said requesting data communication devices; and

50 a link supervisor that transmits from said hub to said requesting access devices, instructions for each requesting access device to transmit at particular times and frequencies, and at particular data rates chosen in accordance with allocations by said bandwidth manager.

10. The digital communications network of claim 9 wherein said common transmission medium comprises a wireless medium.

11. The digital communications network of claim 9 65 wherein said common transmission medium comprises a wired medium.

10

12. The digital communications network of claim 9 wherein said link supervisor comprises a radio link supervisor.

13. In a digital communications network, a hub comprising:

a receiver that receives signals from said data communications devices via said common transmission medium; a bandwidth manager that receives access request messages from requesting data communication devices, said access request messages requesting access to said common transmission medium, and that allocates access to said common transmission medium in both the frequency and time domain among said requesting data communication devices; and a link supervisor that transmits from said hub to said requesting access devices, instructions for each requesting access device to transmit at particular times and frequencies, and at particular data rates chosen in accordance with allocations by said bandwidth manager.

14. The hub of claim 13 wherein said common transmission medium comprises a wired medium.

15. The hub of claim 13 wherein said link supervisor 25 comprises a radio link supervisor.

16. The hub of claim 13 wherein said common transmission medium comprises a wireless medium.

17. The hub of claim 15 wherein said hub further comprises a downstream radio for transmitting signals to said 30 data communication devices.

18. The hub of claim 15 wherein said hub further comprises a plurality of upstream radios for receiving signals from said data communication devices.

19. A data communications device for use in a network, 35 said data communications device comprising:

a bandwidth manager that transmits requests for access to a common transmission medium to a hub; and a link supervisor that receives medium access instructions from said hub, said medium access instructions specifying data rate, transmission time, and transmission frequency for transmissions to said hub, and that controls transmission of information via said common transmission medium in accordance with said specified data rate, transmission time, and transmission frequency.

20. The data communications device of claim 19 wherein said common transmission medium comprises a wired medium.

21. The data communications device of claim 20 wherein said common transmission medium comprises a wireless medium.

22. The data communications device of claim 21 further comprising:

55 at least two downstream radios for communicating at selectable data rates.

23. The data communications device of claim 19 wherein said link supervisor comprises a radio link supervisor.

24. The method of claim 1 further comprising:

assigning a further transmission frame to a data communication device different from said particular data communication device; and

assigning a data rate for said different data communication device to employ in said further transmission frame wherein said particular data communication device and said different data communication device are assigned different data rates.

US 6,463,096 B1

11

25. In a data communication network, apparatus for allocating access to a common transmission medium among a plurality of data communication devices, said apparatus comprising:

means for assigning a transmission frame to a particular data communication device wherein others of said data communication devices do not transmit in said transmission frame;

means for assigning a data rate for said particular data communication device to employ in said transmission frame; and

means for transmitting said transmission frame assignment and said data rate assignment to said particular data communication device.

26. In a data communication network, apparatus for allocating access to a common transmission medium among a plurality of data communication devices, said apparatus comprising:

5
10

12

means for receiving access request messages from requesting data communication devices at a hub, said access request messages requesting access to said common transmission medium;

means for, in response to said access request messages, at said hub, allocating access to said common transmission medium in both the frequency and time domain among said requesting data communication devices; and

means for transmitting from said hub to said requesting access devices, instructions for each requesting access device to transmit at particular times and frequencies, and at particular data rates chosen according to said allocating step.

* * * * *

EXHIBIT 27

PROVISIONAL APPLICATION

Atty. Docket No. 17814-10.00
"Express Mail" Label No. EM 284 724 885US
Date of Deposit November 24, 1997

I hereby certify that this is being deposited with the U.S. Postal Service "Express Mail Post Office to Addressee" service under 37 CFR 1.10 on the date indicated above, addressed to the Asst. Commissioner for Patents, Box Provisional Appln., Washington, DC 20231.

By: John P. Bong

**BOX PROVISIONAL PATENT APPLICATION
ASST. COMMISSIONER FOR PATENTS
Washington, D. C. 20231**

Sir:

Transmitted herewith for filing is a provisional patent application under 37 CFR 1.53(b)(2) of:

LAST NAME	FIRST NAME	MIDDLE INITIAL	RESIDENCE (CITY/STATE/COUNTRY)
Riddle	Guy		Los Gatos, CA, U.S.A.
Packer	Robert	L.	Los Gatos, CA, U.S.A.

Title: METHOD FOR AUTOMATICALLY CLASSIFYING TRAFFIC IN A POLICY BASED BANDWIDTH ALLOCATION SYSTEM

Enclosed are:

30 pages of the specification, claims and abstract.
 8 sheet(s) of informal drawing(s).

A verified statement to establish small entity status under 37 CFR 1.9 and 37 CFR 1.27.

The invention was made by or under a contract with the following agency of the United States Government: _____ under Government contract number: _____

Declaration and Power of Attorney (not signed).
 Appendix A - 31 pages.

We are not paying the fee in this case at this time.

2 extra copies of this sheet are enclosed.

Respectfully submitted,

Correspondence Address:

TOWNSEND and TOWNSEND and CREW LLP

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, CA 94111-3834
Telephone: (650) 326-2400
Fax: (650) 326-2422


Paul A. Durdik
Reg. No.: 37,819
Attorneys for Applicant

Attorney Docket No. 17814-10.00

PROVISIONAL
PATENT APPLICATION

**METHOD FOR AUTOMATICALLY CLASSIFYING TRAFFIC IN A
POLICY BASED BANDWIDTH ALLOCATION SYSTEM**

Inventors:

**Guy Riddle, a United States citizen,
residing at 18243 Knuth Road,
Los Gatos, CA 95033; and**

**Robert L. Packer, a United States citizen,
residing at 16095 Redwood Lodge Road,
Los Gatos, CA 95036.**

Assignee:

**Packeteer, Inc.
(a corporation of Delaware)
307 Orchard City Drive, Suite 305
Campbell, CA 95008**

**TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, 8th Floor
San Francisco, California 94111-3834
(650) 326-2400**

PATENT

Attorney Docket No. 17814-10.00

METHOD FOR AUTOMATICALLY CLASSIFYING TRAFFIC IN A POLICY BASED BANDWIDTH ALLOCATION SYSTEM

COPYRIGHT NOTICE

A portion of the disclosure of this patent document contains material which is subject to copyright protection. The copyright owner has no objection to the facsimile reproduction by anyone of the patent document or the patent disclosure as it appears in the Patent and Trademark Office patent file or records, but otherwise reserves all copyright rights whatsoever.

CROSS-REFERENCE TO RELATED APPLICATIONS

The following related commonly-owned copending U.S. Provisional Patent Application is being filed concurrently and is hereby incorporated by reference in its entirety for all purposes: U.S. Provisional Patent Application Serial No. _____, in the name of Guy Riddle, entitled "Method for Automatically Determining a Traffic Policy in a Policy Based Bandwidth Allocation System," (attorney docket number 17814-9.00), which relates to a determining a default traffic policy.

Further, this application makes reference to the following commonly owned U.S. Patent Applications, which are incorporated by reference herein in their entirety for all purposes:

Copending U.S. Patent Application Serial No. 08/762,828, in the name of Robert L. Packer, entitled "Method for Rapid Data Rate Detection in a Packet Communication Environment Without Data Rate Supervision," relates to a technique for automatically determining the data rate of a TCP connection;

Copending U.S. Patent Application Serial No. _____, in the name of Robert L. Packer, entitled "Method for Managing Flow Bandwidth Utilization at Network, Transport and Application Layers in Store and Forward Network," (attorney docket number 17814-5.10) relates to a technique for automatically allocating bandwidth based upon data rates of TCP connections according to a hierarchical classification paradigm.

Further, this application makes reference to the following U.S. Patent

Application:

Copending U.S. Patent Application Serial No. 08/742,994, in the name of Robert L. Packer, entitled "Method for Explicit Data Rate Control in a Packet Communication Environment Without a Data Rate Supervision," relates to a technique for automatically scheduling TCP packets for transmission.

PAPER APPENDIX

The following paper appendices are included herewith and incorporated by reference in their entirety for all purposes:

Appendix A: Source code listing of automatic classification processing in an embodiment of the invention comprising thirty-one (31) sheets.

BACKGROUND OF THE INVENTION

This invention relates to digital packet telecommunications, and particularly to management of network bandwidth based on information ascertainable from multiple layers of OSI network model. It is particularly useful in conjunction with bandwidth allocation mechanisms employing traffic classification in a digitally-switched packet telecommunications environment normally not subject to data flow rate control.

The ubiquitous TCP/IP protocol suite, which implements the world-wide data communication network environment called the Internet and is also used in private networks (Intranets), intentionally omits explicit supervisory function over the rate of data transport over the various media which comprise the network. While there are certain perceived advantages, this characteristic has the consequence of juxtaposing very high-speed packet flows and very low-speed packet flows in potential conflict for network resources, which results in inefficiencies. Certain pathological loading conditions can result in instability, overloading and data transfer stoppage. Therefore, it is desirable to provide some mechanism to optimize efficiency of data transfer while minimizing the risk of data loss. Early indication of the rate of data flow which can or must be supported is imperative. In fact, data flow rate capacity information is a key factor for use in resource allocation decisions. For example, if a particular path is inadequate to accommodate a high rate of data flow, an alternative route can be sought out.

Internet/Intranet technology is based largely on the TCP/IP protocol suite, where IP, or Internet Protocol, is the network layer protocol and TCP, or Transmission Control Protocol, is the transport layer protocol. At the network level, IP provides a "datagram" delivery service. By contrast, TCP builds a transport level service over the 5 datagram service to provide guaranteed, sequential delivery of a byte stream between two IP hosts.

TCP flow control mechanisms operate exclusively at the end stations to limit the rate at which TCP endpoints emit data. However, TCP lacks explicit data rate control. The basic flow control mechanism is a sliding window, superimposed on a range 10 of bytes beyond the last explicitly-acknowledged byte. Its sliding operation limits the amount of unacknowledged transmissible data that a TCP endpoint can emit.

Another flow control mechanism is a congestion window, which is a refinement of the sliding window scheme, which employs conservative expansion to fully utilize all of the allowable window. A component of this mechanism is sometimes 15 referred to as "slow start".

The sliding window flow control mechanism works in conjunction with the Retransmit Timeout Mechanism (RTO), which is a timeout to prompt a retransmission of unacknowledged data. The timeout length is based on a running average of the Round 20 Trip Time (RTT) for acknowledgment receipt, i.e. if an acknowledgment is not received within (typically) the smoothed RTT + 4*mean deviation, then packet loss is inferred and the data pending acknowledgment is retransmitted.

Data rate flow control mechanisms which are operative end-to-end without explicit data rate control draw a strong inference of congestion from packet loss (inferred, typically, by RTO). TCP end systems, for example, will 'back-off', i.e., inhibit 25 transmission in increasing multiples of the base RTT average as a reaction to consecutive packet loss.

Bandwidth Management in TCP/IP Networks

Conventional bandwidth management in TCP/IP networks is accomplished 30 by a combination of TCP end systems and routers which queue packets and discard packets when certain congestion thresholds are exceeded. The discarded, and therefore unacknowledged, packet serves as a feedback mechanism to the TCP transmitter. (TCP

end systems are clients or servers running the TCP transport protocol, typically as part of their operating system.)

The term "bandwidth management" is often used to refer to link level bandwidth management, e.g. multiple line support for Point to Point Protocol (PPP).

5 Link level bandwidth management is essentially the process of keeping track of all traffic and deciding whether an additional dial line or ISDN channel should be opened or an extraneous one closed. The field of this invention is concerned with *network* level bandwidth management, i.e. policies to assign available bandwidth from a single logical link to network flows.

10 In a copending U.S. Patent Application Serial No. 08/742,994, in the name of Robert L. Packer, entitled "Method for Explicit Data Rate Control in a Packet Communication Environment Without Data Rate Supervision," a technique for automatically scheduling TCP packets for transmission is disclosed. Furthermore, in a copending U.S. Patent Application Serial No. 08/762,828, in the name of Robert L. Packer, entitled "Method for Rapid Data Rate Detection in a Packet Communication Environment Without Data Rate Supervision," a technique for automatically determining the data rate of a TCP connection is disclosed. Furthermore, in a copending U.S. Patent Application Serial No. _____, in the name of Robert L. Packer, entitled "Method for Managing Flow Bandwidth Utilization at Network, Transport and Application Layers in Store and Forward Network," (attorney docket number 17814-5.10) a technique for automatically allocating bandwidth based upon data rates of TCP connections according to a hierarchical classification paradigm is disclosed.

15 Automated tools assist the network manager in configuring and managing the network equipped with the rate control techniques described in these copending applications. In a related copending application, a tool is described which enables a network manager to automatically produce policies for traffic being automatically detected in a network. It is described in a copending U.S. Provisional Patent Application Serial No. _____, in the name of Guy Riddle, entitled "Method for Automatically Determining a Traffic Policy in a Policy Based Bandwidth Allocation System," (attorney docket number 17814-9.00). The subject of the present invention is also a tool designed 20 to assist the network manager.

While these efforts teach methods for solving problems associated with scheduling transmissions, automatically determining data flow rate on a TCP connection, allocating bandwidth based upon a classification of network traffic and automatically determining a policy, respectively, there is no teaching in the prior art of methods for 5 automatically classifying packet traffic based upon information gathered from a multiple layers in a multi-layer protocol network.

Bandwidth has become the expensive commodity of the '90s, as traffic expands faster than resources, the need to "prioritize" a scarce resource, becomes ever more critical. One way to solve this is by applying "policies" to control traffic classified 10 as to type of service required in order to more efficiently match resources with traffic.

Traffic may be classified by type, e.g. E-mail, web surfing, file transfer, at various levels. For example, to classify by network paradigm, examining messages for an IEEE source/destination service access point (SAP) or a sub-layer access protocol (SNAP) yields a very broad indicator, i.e., SNA or IP. More specific types exist, such as whether 15 an IP protocol field in an IP header indicates TCP or UDP. Well known connection ports provide indications at the application layer, i.e., SMTP or HTTP.

Classification is not new. Firewall products like "CheckPoint FireWall-1," a product of CheckPoint Software Technologies, Inc., a company with headquarters in Redwood City, CA., have rules for matching traffic. Bandwidth managers such as 20 "Aponet," a product of Aponet, Inc., a company with headquarters in San Jose, CA., classify by destination. The PacketShaper, a product of Packeteer, Inc., a company with headquarters in Campbell, CA., allows a user to manually enter rules to match various traffic types for statistical tracking, i.e., counting by transaction, byte count, rates, etc. However, manual rule entry requires a level of expertise that limits the appeal for such a 25 system to network savvy customers. What is really needed is a method for analyzing real traffic in a customer's network and automatically producing a list of the "found traffic."

SUMMARY OF THE INVENTION

According to the invention, in a packet communication environment, a 30 method is provided for automatically classifying packet flows for use in allocating bandwidth resources by a rule of assignment of a service level. The method comprises applying individual instances of traffic classification paradigms to packet network flows based on selectable information obtained from a plurality of layers of a multi-layered

communication protocol in order to define a characteristic class, then mapping the flow to the defined traffic class. It is useful to note that the automatic classification is sufficiently robust to classify a complete enumeration of the possible traffic.

5 An advantage of traffic classification techniques according to the present invention is that network managers need not know the technical aspects of each kind of traffic in order to configure traffic classes.

A further advantage of the present invention is that traffic classes may include information such as a URI for web traffic.

10 A yet further advantage of the present invention is that service aggregates bundle traffic to provide a convenience to the user, by clarifying processing and enables the user to obtain group counts of all parts comprising a service.

The invention will be better understood upon reference to the following detailed description in connection with the accompanying drawings.

15 BRIEF DESCRIPTION OF THE DRAWINGS

Fig. 1A depicts a representative client server relationship in accordance with a particular embodiment of the invention;

Fig. 1B depicts a functional perspective of the representative client server relationship in accordance with a particular embodiment of the invention;

20 Fig. 1C depicts a representative internetworking environment in accordance with a particular embodiment of the invention;

Fig. 1D depicts a relationship diagram of the layers of the TCP/IP protocol suite;

Figs. 2A-2B depict representative divisions of bandwidth;

25 Fig. 3 depicts a component diagram of processes and data structures in accordance with a particular embodiment of the invention; and

Figs. 4A-4B depict flowcharts of process steps in automatically classifying traffic in accordance with a particular embodiment of the invention.

DETAILED DESCRIPTION OF SPECIFIC EMBODIMENTS

A preferable embodiment of a flow bandwidth management system according to the invention has been reduced to practice and will be made available under the trade name "PacketShaperTM."

5

1.0 Introduction

The present invention provides techniques to automatically classify a plurality of heterogeneous packets in a packet telecommunications system for management of network bandwidth in systems such as a private area network, a wide area network or an internetwork. Systems according to the present invention enable network managers to: automatically define traffic classes, for which policies may then be created for specifying service levels for the traffic classes and isolating bandwidth resources associated with certain traffic classes. Inbound as well as outbound traffic may be managed. Table 1 provides a definitional list of terminology used herein.

15

LIST OF DEFINITIONAL TERMS

ADMISSIONS CONTROL A policy invoked whenever a system according to the invention detects that a guaranteed information rate cannot be maintained. An admissions control policy is analogous to a busy signal in the telephone world.

CLASS SEARCH ORDER A search method based upon traversal of a N-ary tree data structure containing classes.

25

COMMITTED INFORMATION

RATE

(CIR)

A rate of data flow allocated to reserved service traffic for rate based bandwidth allocation for a committed bandwidth. Also called a guaranteed information rate (GIR).

30

EXCEPTION

A class of traffic provided by the user which supersedes an automatically determined classification order.

EXCESS INFORMATION
RATE
(EIR)

A rate of data flow allocated to reserved service traffic for rate based bandwidth allocation for uncommitted bandwidth resources.

5

FLOW

A flow is a single instance of a traffic class. For example, all packets in a TCP connection belong to the same flow. As do all packets in a UDP session.

10

GUARANTEED
INFORMATION RATE
(GIR)

A rate of data flow allocated to reserved service traffic for rate based bandwidth allocation for a committed bandwidth. Also called a committed information rate (CIR).

15

HARD ISOLATION

Hard isolation results from the creation of an entirely separate logical channel for a designated set of classes.

20

INSIDE

On the system side of an access link. Outside clients and servers are on the other side of the access link.

ISOLATION

Isolation is the degree that bandwidth resources are allocable to traffic classes.

25

OUTSIDE

On the opposite side of an access link as viewed from the perspective of the system on which the software resides.

PARTITION

Partition is an arbitrary unit of network resources.

30

POLICY

A rule for the assignment of a service level to a flow.

POLICY INHERITANCE

A method for assigning policies to flows for which no policy exists in a hierarchical arrangement of policies. For

example, if a flow is determined to be comprised of FTP packets for Host A, and no corresponding policy exists, a policy associated with a parent node, such as an FTP policy, may be located and used. See also POLICY SEARCH ORDER.

5

POLICY BASED SCALING

10

An adjustment of a requested data rate for a particular flow based upon the policy associated with the flow and information about the flow's potential rate.

RESERVED SERVICE

15

Reserved service is a service level intended for traffic which "bursts" or sends chunks of data. Reserved service is defined in terms of a scaled rate.

SCALED RATE

20

Assignment of a data rate based upon detected speed.

SERVICE LEVEL

A service paradigm having a combination of characteristics defined by a network manager to handle a particular class of traffic. Service levels may be designated as either reserved or unreserved.

SOFT ISOLATION

25

Restricting GIR allocated for traffic classes in a partition.

TARGET RATE

30

A target rate is a combination of a guaranteed rate and an excess rate. Target rate is a policy-based paradigm. Excess rate is allocated by systems according to the invention from bandwidth that is not consumed by reserved service. Policies will demand excess rate at a given priority and systems according to the invention satisfy this demand by a priority level.

TRAFFIC CLASS	All traffic between a client and a server endpoints. A single instance of a traffic class is called a flow. Traffic classes have properties or class attributes such as, <i>directionality</i> , which is the property of traffic to be flowing inbound or outbound.
5	UNRESERVED SERVICE Unreserved service is a service level defined in terms of priority in which no reservation of bandwidth is made.

10

Table 1

1.1 Hardware Overview

The method for automatically classifying heterogeneous packets in a packet telecommunications environment of the present invention is implemented in the C programming language and is operational on a computer system such as shown in Fig. 1A. This invention may be implemented in a client-server environment, but a client-server environment is not essential. This figure shows a conventional client-server computer system which includes a server 20 and numerous clients, one of which is shown as client 25. The use of the term "server" is used in the context of the invention, wherein the server receives queries from (typically remote) clients, does substantially all the processing necessary to formulate responses to the queries, and provides these responses to the clients. However, server 20 may itself act in the capacity of a client when it accesses remote databases located at another node acting as a database server.

The hardware configurations are in general standard and will be described only briefly. In accordance with known practice, server 20 includes one or more processors 30 which communicate with a number of peripheral devices via a bus subsystem 32. These peripheral devices typically include a storage subsystem 35, comprised of a memory subsystem 35a and a file storage subsystem 35b holding computer programs (e.g., code or instructions) and data, a set of user interface input and output devices 37, and an interface to outside networks, which may employ Ethernet, Token Ring, ATM, IEEE 802.3, ITU X.25, Serial Link Internet Protocol (SLIP) or the public switched telephone network. This interface is shown schematically as a "Network

Interface" block 40. It is coupled to corresponding interface devices in client computers via a network connection 45.

Client 25 has the same general configuration, although typically with less storage and processing capability. Thus, while the client computer could be a terminal or 5 a low-end personal computer, the server computer is generally a high-end workstation or mainframe, such as a SUN SPARC server. Corresponding elements and subsystems in the client computer are shown with corresponding, but primed, reference numerals.

Bus subsystem 32 is shown schematically as a single bus, but a typical system has a number of buses such as a local bus and one or more expansion buses (e.g., 10 ADB, SCSI, ISA, EISA, MCA, NuBus, or PCI), as well as serial and parallel ports.

Network connections are usually established through a device such as a network adapter on one of these expansion buses or a modem on a serial port. The client computer may be a desktop system or a portable system.

15 The user interacts with the system using interface devices 37' (or devices 37 in a standalone system). For example, client queries are entered via a keyboard, communicated to client processor 30', and thence to modem or network interface 40' over bus subsystem 32'. The query is then communicated to server 20 via network connection 45. Similarly, results of the query are communicated from the server to the client via 20 network connection 45 for output on one of devices 37' (say a display or a printer), or may be stored on storage subsystem 35'.

Fig. 1B is a functional diagram of a computer system such as that of Fig. 1A. Fig. 1B depicts a server 20, and a representative client 25 of a plurality of clients which may interact with the server 20 via the Internet 45 or any other communications method. Blocks to the right of the server are indicative of the processing steps and 25 functions which occur in the server's program and data storage indicated by blocks 35a and 35b in Fig. 1A. A TCP/IP "stack" 44 works in conjunction with Operating System 42 to communicate with processes over a network or serial connection attaching Server 20 to Internet 45. Web server software 46 executes concurrently and cooperatively with other processes in server 20 to make data objects 50 and 51 available to requesting 30 clients. A Common Gateway Interface (CGI) script 55 enables information from user clients to be acted upon by web server 46, or other processes within server 20. Responses to client queries may be returned to the clients in the form of a Hypertext Markup

Language (HTML) document outputs which are then communicated via Internet 45 back to the user.

Client 25 in Fig. 1B possesses software implementing functional processes operatively disposed in its program and data storage as indicated by block 35a' in Fig.

5 1A. TCP/IP stack 44', works in conjunction with Operating System 42' to communicate with processes over a network or serial connection attaching Client 25 to Internet 45.

Software implementing the function of a web browser 46' executes concurrently and cooperatively with other processes in client 25 to make requests of server 20 for data objects 50 and 51. The user of the client may interact via the web browser 46' to make 10 such queries of the server 20 via Internet 45 and to view responses from the server 20 via Internet 45 on the web browser 46'.

Network Overview

Fig. 1C is illustrative of the internetworking of a plurality of clients such as client 25 of Figs. 1A and 1B and a plurality of servers such as server 20 of Figs. 1A and 1B as described herein above. In Fig. 1C, network 70 is an example of a Token Ring or frame oriented network. Network 70 links host 71, such as an IBM RS6000 RISC workstation, which may be running the AIX operating system, to host 72, which is a personal computer, which may be running Windows 95, IBM OS/2 or a DOS operating system, and host 73, which may be an IBM AS/400 computer, which may be running the OS/400 operating system. Network 70 is internetworked to network 60 via a system gateway which is depicted here as router 75, but which may also be a gateway having a firewall or a network bridge. Network 60 is an example of an Ethernet network that interconnects host 61, which is a SPARC workstation, which may be running SUNOS 20 operating system with host 62, which may be a Digital Equipment VAX6000 computer which may be running the VMS operating system.

Router 75 is a network access point (NAP) of network 70 and network 60. Router 75 employs a Token Ring adapter and Ethernet adapter. This enables router 75 to interface with the two heterogeneous networks. Router 75 is also aware of the Inter-network Protocols, such as ICMP ARP and RIP, which are described herein below.

30 Fig. 1D is illustrative of the constituents of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite. The base layer of the TCP/IP protocol suite is the physical layer 80, which defines the mechanical, electrical, functional and

procedural standards for the physical transmission of data over communications media, such as, for example, the network connection 45 of Fig. 1A. The physical layer may comprise electrical, mechanical or functional standards such as whether a network is packet switching or frame-switching; or whether a network is based on a Carrier Sense 5 Multiple Access/Collision Detection (CSMA/CD) or a frame relay paradigm.

Overlying the physical layer is the data link layer 82. The data link layer provides the function and protocols to transfer data between network resources and to detect errors that may occur at the physical layer. Operating modes at the datalink layer comprise such standardized network topologies as IEEE 802.3 Ethernet, IEEE 802.5 10 Token Ring, ITU X.25, or serial (SLIP) protocols.

Network layer protocols 84 overlay the datalink layer and provide the means for establishing connections between networks. The standards of network layer protocols provide operational control procedures for internetworking communications and routing information through multiple heterogenous networks. Examples of network layer protocols are the Internet Protocol (IP) and the Internet Control Message Protocol (ICMP). The Address Resolution Protocol (ARP) is used to correlate an Internet address and a Media Access Address (MAC) for a particular host. The Routing Information Protocol (RIP) is a dynamic routing protocol for passing routing information between hosts on networks. The Internet Control Message Protocol (ICMP) is an internal protocol 15 for passing control messages between hosts on various networks. ICMP messages provide feedback about events in the network environment or can help determine if a path exists to a particular host in the network environment. The latter is called a "Ping". The Internet Protocol (IP) provides the basic mechanism for routing packets of information in the Internet. IP is a non-reliable communication protocol. It provides a "best efforts" 20 delivery service and does not commit network resources to a particular transaction, nor does it perform retransmissions or give acknowledgments.

The transport layer protocols 86 provide end-to-end transport services across multiple heterogenous networks. The User Datagram Protocol (UDP) provides a connectionless, datagram oriented service which provides a non-reliable delivery 25 mechanism for streams of information. The Transmission Control Protocol (TCP) provides a reliable session-based service for delivery of sequenced packets of information across the Internet. TCP provides a connection oriented reliable mechanism for information delivery.

The session, or application layer 88 provides a list of network applications and utilities, a few of which are illustrated here. For example, File Transfer Protocol (FTP) is a standard TCP/IP protocol for transferring files from one machine to another. FTP clients establish sessions through TCP connections with FTP servers in order to 5 obtain files. Telnet is a standard TCP/IP protocol for remote terminal connection. A Telnet client acts as a terminal emulator and establishes a connection using TCP as the transport mechanism with a Telnet server. The Simple Network Management Protocol (SNMP) is a standard for managing TCP/IP networks. SNMP tasks, called "agents", monitor network status parameters and transmit these status parameters to SNMP tasks 10 called "managers." Managers track the status of associated networks. A Remote Procedure Call (RPC) is a programming interface which enables programs to invoke remote functions on server machines. The Hypertext Transfer Protocol (HTTP) facilitates the transfer of data objects across networks via a system of uniform resource indicators (URI).

15 The Hypertext Transfer Protocol is a simple protocol built on top of Transmission Control Protocol (TCP). It is the mechanism which underlies the function of the World Wide Web. The HTTP provides a method for users to obtain data objects from various hosts acting as servers on the Internet. User requests for data objects are made by means of an HTTP request, such as a GET request. A GET request as depicted 20 below is comprised of 1) the GET request keyword; followed by 2) the full path of the data object; followed by 3) the name of the data object; followed by 4) an HTTP protocol version, such as "HTTP/1.0". In the GET request shown below, a request is being made for the data object with a path name of "/pub/" and a name of "MyData.html":

25 GET /pub/MyData.html HTTP-Version (1)

30 Processing of a GET request entails the establishing of an TCP/IP connection with the server named in the GET request and receipt from the server of the data object specified. After receiving and interpreting a request message, a server responds in the form of an HTTP RESPONSE message.

Response messages begin with a status line comprising a protocol version followed by a numeric Status Code and an associated textual Reason Phrase. These

elements are separated by space characters. The format of a status line is depicted in line (2):

Status-Line = HTTP-Version Status-Code Reason-Phrase (2)

5

The status line always begins with a protocol version and status code, e.g., "HTTP/1.0 200 ". The status code element is a three digit integer result code of the attempt to understand and satisfy a prior request message. The reason phrase is intended to give a short textual description of the status code.

10 The first digit of the status code defines the class of response. There are five categories for the first digit. 1XX is an information response. It is not currently used. 2XX is a successful response, indicating that the action was successfully received, understood and accepted. 3XX is a redirection response, indicating that further action must be taken in order to complete the request. 4XX is a client error response. This indicates a bad syntax in the request. Finally, 5XX is a server error. This indicates that the server failed to fulfill an apparently valid request.

15

2.0 Traffic Class

20 A traffic class is broadly defined as traffic between one or more clients and one or more servers. A single instance of a traffic class is called a flow. Traffic classes have the property, or class attribute, of being directional, i.e. all traffic flowing inbound will belong to different traffic classes and be managed separately from traffic flowing outbound. The directional property enables asymmetric classification and control of traffic, i.e., inbound and outbound flows belong to different classes which may be 25 managed independent of one another.

25 Traffic classes may be defined at any level of the TCP/IP protocol. For example, at the IP level, traffic may be defined as only those flows between a set of inside and outside IP addresses or domain names. An example of such a low level traffic class definition would be all traffic between my network and other corporate offices 30 throughout the Internet. At the application level, traffic classes may be defined for specific URIs within a web server. Traffic classes may be defined having "Web aware" class attributes. For example, a traffic class could be created such as all URIs matching "*.html" for all servers, or all URIs matching "*.gif" for server X, or for access to server

Y with URI "/sales/*" from client Z, wherein '*' is a wildcard character, i.e., a character which matches all other character combinations. Traffic class attributes left unspecified will simply match any value for that attribute. For example, a traffic class that accesses data objects within a certain directory path of a web server is specified by a URI of the 5 directory path to be managed, e.g. "/sales/*".

2.1 Classifying Traffic

The present invention provides a method for classifying traffic according to a definable set of classification attributes selectable by the manager, including selecting a 10 subset of traffic of interest to be classified. The invention provides the ability to classify and search traffic based upon multiple orthogonal classification attributes.

Traffic class membership may be hierarchical. Thus, a flow may be classified by a series 15 of steps through a traffic class tree, with the last step (i.e., at the leaves on the classification tree) mapping the flow to a policy. The policy is a rule of assignment for flows. For example, the first step in classification may be to classify a flow as web traffic, the next may further classify this flow as belonging to server X, and the final classification may be a policy for URI "*.avi".

A classification tree is a data structure representing the hierarchical aspect 20 of traffic class relationships. Each node of the classification tree represents a class, and has a traffic specification, i.e., a set of attributes or characteristics describing the traffic, and a mask associated with it. Leaf nodes of the classification tree may contain policies. According to a particular embodiment, the classification process checks at each level if 25 the flow being classified matches the attributes of a given traffic class. If it does, processing continues down to the links associated with that node in the tree. If it does not, the class at the level that matches determines the policy for the flow being classified. If no policy specific match is found, the flow is assigned the default policy.

In a preferable embodiment, the classification tree is an N-ary tree with its 30 nodes ordered by specificity. For example, in classifying a particular flow in a classification tree ordered first by organizational departments, the attributes of the flow are compared with the traffic specification in each successive department node and if no match is found, then processing proceeds to the next subsequent department node. If no match is found, then the final compare is a default "match all" category. If, however, a match is found, then classification moves to the children of this department node. The

child nodes may be ordered by an orthogonal paradigm such as, for example, "service type." Matching proceeds according to the order of specificity in the child nodes. Processing proceeds in this manner, traversing downward and from left to right in Figs. 2A and 2B, which describe a classification tree, searching the plurality of orthogonal paradigms. Key to implementing this a hierarchy is that the nodes are arranged in decreasing order of specificity. This permits search to find the most specific class for the traffic before more general.

Table 2 depicts components from which Traffic classes may be built. Note that the orientation of the server (inside or outside) is specified. And as noted above, any traffic class component may be unspecified, i.e. set to match any value.

Traffic Class Components	
Client Side	Server Side
IP Address/Domain Name	IP Address/Domain Name
	TCP or UDP Service, e.g. WWW, FTP, RealAudio, etc. URI for Web Service, e.g. "*.html", "*gif", "/sales/*", etc.

25 Table 2

Figs. 2A and 2B depict representative allocations of bandwidth made by a hypothetical network manager as an example. In Fig. 2A, the network manager has decided to divide her network resources first by allocating bandwidth between Departments A and B. Fig 2A shows the resulting classification tree, in which Department A bandwidth resources 202 and Department B bandwidth resources 204 each have their own nodes representing a specific traffic class for that department. Each traffic class may have a policy attribute associated with it. For example, in Fig. 2A, the Department A resources node 202 has the policy attribute Inside Host Subnet A associated

with it. Next, the network manager has chosen to divide the bandwidth resources of Department A among two applications. She allocates an FTP traffic class 206 and a World Wide Web server traffic class 208. Each of these nodes may have a separate policy attribute associated with them. For example, in Fig. 2A, the FTP node 206 for has 5 an attribute Outside port 20 associated with it. Similarly, the network manager has chosen to divide network bandwidth resources of Department B into an FTP server traffic class 210 and a World Wide Web server traffic class 212. Each may have their own respective policies.

Fig. 2B shows a second example, wherein the network manager has 10 chosen to first divide network bandwidth resource between web traffic and TCP traffic. She creates three traffic nodes, a web traffic node 220, a TCP traffic node 224 and a default node 225. Next, she divides the web traffic among two organizational departments by creating a Department A node 226, and a Department B node 228. Each 15 may have its own associated policy. Similarly, she divides TCP network bandwidth into separate traffic classes by creating a Department A node 230 and a Department B node 232. Each represents a separate traffic class which may have its own policy.

All traffic which does not match any user specified traffic class falls into 20 an automatically created default traffic class which has a default policy. In Fig 2A, the default category is depicted by a default node 205, and in Fig. 2B, the default category is depicted by a default node 225.

3.0 Automatically Classifying Traffic

3.1 Automatic Traffic Classification

Network traffic is automatically classified under existing classes, 25 beginning with the broadest classes, an inbound traffic class and an outbound traffic class, in protocol layer independent categories. For example, a particular instance of traffic may be classified according to its transport layer characteristics, e.g., Internet Protocol port number, as well as its application layer information, e.g., SMTP. Characteristics such as MIME types may also be automatically identified. Standard protocols, such as, IPX, 30 SNA, and services, such as, SMTP and FTP are recognized for automatic classification. Classification is performed to the most specific level determinable. For example, in select embodiments, non-IP traffic, such as SNA, may be classified only by protocol, whereas Internet Protocol traffic may be classified to the /etc/services level. Classification beyond

a terminal classification level is detected and prevented. For example, in a select embodiment, a class matching "ipx" or "nntp" will not be further automatically classified.

3.1.1 Service Aggregates

5 A service aggregate is provided for certain applications that use more than one connection in a particular conversation between a client and a server. For example, an FTP client in conversation with an FTP server employs a command channel and a transfer channel, which are distinct TCP sessions on two different ports. In cases where 10 two or three TCP or UDP sessions exist for each conversation between one client and one server, it is useful to provide a common traffic class i.e., the service aggregate, containing the separate conversations. In practice, these types of conversations are between the same two hosts, but use different ports. According to the invention, a class is created with a plurality of traffic specifications, each matching various component conversations.

15 3.1.2 Subclassification Under Specified Criterion

Subclassification of traffic into a tree is performed by matching the hosts and then searching for particular services. Traffic specifications are aggregate kinds of traffic for a traffic class, e.g., different components of FTP may reside under class FTP. Subclassification is performed by first locating a class that matches, and then performing 20 finer grade matchings. Traffic specifications have a default for each traffic class node. Processing commences with a decision on what traffic is to be subclassified. A marker is placed in the match_all default node so that when match processing reaches the marker, the autoclassification processing depicted in flowchart 403, determines that it has not found an existing class for the traffic being classified.

25 3.1.3 Default Suggested Policies

A default policy may be suggested or, in select embodiments, automatically applied, to a traffic class which has been automatically classified. Applying suggested or default policies for a new class at a user's option is described in a 30 copending, commonly owned, U.S. Provisional Patent Application Serial No. _____, entitled, "Method for Automatically Determining a Traffic Policy in a Policy Based Bandwidth Allocation System," (attorney docket number 17814-9.00), which is incorporated herein by reference in its entirety for all purposes.

3.1.4 Analysis of Data in Determining Traffic Class

In a preferable embodiment, classification can extend to examination of the data contained in a flow's packets. Certain traffic may be distinguished by a signature even if it originates with a server run on a non-standard port, for example, an HTTP conversation on port 8080 would not be otherwise determinable as HTTP from the port number. Further analysis of the data is conducted in order to determine classification in instances where: 1) FTP commands are used to define server ports, 2) HTTP protocol is used for non-web purposes. The data is examined for indication of push traffic, such as pointcast, which uses HTTP as a transport mechanism. These uses may be isolated and classified into a separate class. Marimba and pointcast can be distinguished by looking into the data for a signature content header in the get request. Pointcast has URLs that begin with "/FIDO-1/." Other applications in which protocol can be inferred from data include Telnet traffic. Both tn3270 and tn3270E (emulation) may be detected by looking into data and given a different class. Telnet traffic has option negotiations which may indicate an appropriate class.

3.1.5 Identity of Traffic Based Upon Resource Creator's Class

A traffic class may be inferred from determining the identity of the creator of a resource used by the traffic class. For example, the identity of traffic using a certain connection can be determined by finding the identity of the creator of the connection. This method is used to detect Real Time Protocol (RTP) for point-to-point telephony, RTP for broadcast streaming, CCITT/ITU H320- telephony over ISDN, H323- internet telephony over the internet (bidirectional) and RTSP real time streaming protocol for movies (unidirectional).

3.1.6 Dynamic Ports:

Applications having a well known port for a server may make use of dynamic ports. Some applications will send initial messages across a first connection, then negotiate a dynamic port for further conversation. During the existence of a connection, both endpoints are known. A check is made for two simultaneous connections to the same, non well-known port, at same time from different locations. This condition is indicative of a connection port for some application. Varieties of the dynamic port exist in applications. Certain dynamic ports are incorporated into a client.